Methodological Advances in Using Negative Control Endpoints to Improve Vaccine Effectiveness Estimation

Ethan Ashby

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Bo Zhang, Chair

Ting Ye, Chair

Holly Janes

Helen Chu

Program Authorized to Offer Degree:
University of Washington Department of Biostatistics

University of Washington

Abstract

Methodological Advances in Using Negative Control Endpoints to Improve Vaccine Effectiveness Estimation

Ethan Ashby

Co-Chairs of the Supervisory Committee:
Bo Zhang
Assistant Professor, Biostatistics, Bioinformatics and Epidemiology Program, Vaccine and Infectious Disease Division, Fred Hutchinson Cancer Research Center

Ting Ye
Associate Professor, Department of Biostatistics, University of Washington

This dissertation addresses two problems at the frontier of vaccine effectiveness evaluation that are important to vaccine science and public health: (i) obtaining reliable evidence of waning vaccine effectiveness over time and (ii) improving the precision of vaccine efficacy estimates in randomized vaccine trials. To address these challenges, we develop novel statistical methods that leverage negative control outcomes—outcomes unaffected by vaccination yet correlated with the primary outcome through unmeasured common causes. In Part 1, we examine the unmeasured biases affecting commonly used hazard-based estimators of time-varying vaccine effectiveness and propose statistical frameworks for diagnosing and correcting for unmeasured bias using negative control infection endpoints. As a case study, we apply our methods to the Coronavirus Vaccine Efficacy (COVE) study, which led to qualitative reassessment of mRNA-1273 booster efficacy in this setting. In Part 2, we extend statistical methods for covariate adjustment to incorporate valid negative control endpoints in early-and-late-phase randomized vaccine trials. With particular emphasis on HIV-1 vaccine trials, we show that incorporating negative control immune responses and infection outcomes can enhance the efficiency of vaccine effect estimates when negative control and primary outcomes share unmeasured prognostic factors.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF INCLUDED CODE

ACKNOWLEDGMENTS

I would like to thank my dissertation advisors, Bo Zhang (Vaccine and Infectious Disease Division (VIDD), Fred Hutchinson Cancer Research Center (FHCRC)) and Ting Ye (Department of Biostatistics, University of Washington) for their support and mentorship throughout my PhD. I have witnessed how you both identify important scientific problems and address them with rigorous statistical insight – these are traits that I will aspire to throughout my career. I have learned very much from you, it has been a pleasure working with you, and I look forward to continued collaboration opportunities in the future.

I would also like to extend my gratitude to Holly Janes (VIDD, FHCRC) for her mentorship throughout my graduate school experience. I credit Holly with introducing me to the vibrant research environment in VIDD at Fred Hutch, which has kindled a lifelong passion for vaccine research and biostatistical methods development. Holly's leadership in VIDD has been an inspiration to witness, and I thank her for her support of me, those in the VIDD group, and vaccine researchers across the world.

I would also like to thank Peter Gilbert (VIDD, FHCRC), Hana M. El Sahly (Department of Molecular Viorlogy and Microbiology and Medicine, Baylor College of Medicine), and Lindsey R. Baden (Brigham and Womens' Hospital, Harvard Medical School) for their collaboration and leadership on the Coronavirus Vaccine Efficacy (COVE) study, which motivated a significant portion of my dissertation work. I would also like to thank the many other COVE study contributors and participants who helped make my research possible. I would like to especially thank Dean Follmann (formerly Biostatistics Research Branch, National Institute of Allergies and Infectious Diseases) for generously sharing his time and expertise with a junior researcher like me. I would also like to thank Tanya Garcia (Department of Biostatistics, University of North Carolina, Chapel Hill) and Zhewei Zhang (Department of Statistics, Penn State University) for their enthusiastic collaboration on my final dissertation project.

Many thanks as well to the Brown-Janes-Gao lab group for providing a welcoming and constructive environment to share research ideas and solicit feedback on projects. In particular, I would like to thank Elizabeth Brown, Fei Gao, Aaron Hudson, Tracy Dong, Laina Mercer, Angela Dahl (all VIDD, FHCRC) and Marlena Bannick (Seattle Childen's Therapeutics) for helpful conversations that have informed my work.

I would also like to thank my PhD cohort – Ellen Graham, James Peng, Albert Osom, Katie Paulson, Yinxiang Wu, Sungtaek Son, and Jaewon Lim – for their support through my graduate school journey. I would also like to thank a few students outside my cohort – Ameer Dharamshi, Ronan Perry, Rui Wang – whom I the pleasure of taking the core statistical courses with during my time at UW. I would like to thank all the students and faculty of the University of Washington Biostatistics Department, who helped foster a inclusive and exciting research environment that I have treasured being a part of the last five years.

Many thanks to my community both within and outside Seattle for always providing perspective and support throughout my time at UW. You helped me stay positive, evolve as person outside of work, and build memories that I will cherish for a lifetime.

Finally, thank you to my family – my father Mark, my mother Mary Jo, my brother Emerson, and my girlfriend Paula – for their love, support, and encouragement throughout my graduate school journey. You all inspire me, lift me up when I am down, and give me the confidence to tackle challenging and meaningful problems.

DEDICATION

I dedicate this dissertation to my family and good-faith vaccine researchers everywhere.

Chapter 1

INTRODUCTION

1.1 Negative controls and their utility to vaccine science

Vaccines are one of the most effective public health interventions for preventing morbidity and mortality caused by infectious diseases. Immunizations prevent between 3.5 and 5 million deaths annually caused by pathogens.[1] Since the advent of World Health Organization's Expanded Programme on Immunization in 1974, it is estimated that expanded immunization coverage has averted 154 million deaths globally, including 146 million among children less than 5 years old and 101 million infants less than 1 year old.[2] The immense value of vaccines to public health was made evident during the global COVID-19 pandemic. Several Phase III, randomized, double-blinded, placebo-controlled trials of COVID-19 vaccines demonstrated a high level of efficacy against COVID-19 disease and severe disease.[3–8] As a result, vaccines were distributed to the public, which is estimated to have saved millions of lives[9] and enabled the relaxation of pandemic-related restrictions.

The immense public health impact of vaccines underscores the importance of rigorously evaluating their safety and protective effects. Randomized, double-blinded, placebo-controlled vaccine trials therefore play an essential role, as they provide a means to reliably evaluate vaccine efficacy and safety in humans.[10] Under proper conduct, randomized trials enable causal interpretation of observed differences in disease risk between vaccine and placebo recipients.[11,12] Randomization ensures vaccinated and placebo recipients are exchangeable (or balanced) with respect to both measured and unmeasured characteristics that may impact their infection risk. Consequently, differences in disease risk between study arms can be directly attributed to the vaccine rather than to systematic differences between participants. Participant and caregiver blinding further protects against distortion of study results by differences in treatment delivery, outcome assessment, expectation of treatment benefit, or participant behavior.[13]

Despite their central role, randomized vaccine trials cannot address many important scientific and public health questions regarding vaccines. In particular, vaccine effectiveness, meaning the impact of vaccination under real-world conditions, has clear policy implications and must frequently

be evaluated using observational data. Observational data play a crucial role in understanding vaccine effectiveness against emerging viral variants, effectiveness in populations frequently underrepresented in clinical trials (e.g., immunocompromised individuals, pregnant persons, and children), efficacy against post-infection endpoints such as transmission or disease progression, optimal dosing intervals, waning immunity, and the effectiveness of heterologous booster strategies.[14–17]

However, observational studies (which lack a randomized treatment assignment) can produce misleading results due to systematic differences between vaccinated and unvaccinated individuals. A canonical example is how influenza vaccination was associated with marked reduction in pneumonia/influenza hospitalization and all-cause mortality in elderly populations observed across several observational studies.[18] These results were inconsistent with other ecological data, suggesting that uncontrolled confounding through health-seeking behavior, baseline health status, or exposure may have inflated the magnitude of the association.[19] Another example is the null or negative association between COVID-19 vaccination and COVID-19 outcomes observed in several observational studies, which is hypothesized to be the result of unmeasured confounding or selection bias.[20]

Therefore, strategies to detect and adjust for unmeasured confounding in observational vaccine studies are needed. Borrowing ideas from laboratory science, negative controls have become increasingly popular tools in causal inference and epidemiology for detecting and correcting for bias in non-randomized analyses.[21–25] Negative controls can be divided into two classes: negative control outcomes (NCOs) and negative control exposures (NCEs). A NCO is a variable known not to be causally affected by the exposure of interest, while a NCE is a variable known not to causally affect the outcome of interest. Since there is no causal relationship between the primary outcome/exposure and negative controls, observed associations involving negative controls are diagnostic evidence of residual confounding or other sources of bias. When paired with additional assumptions and appropriate statistical techniques, negative controls can be used to reduce or eliminate bias from statistical parameters.[26–29] Negative controls have been advocated as a general bias detection tools, including accounting for measurement error and selection bias[30] and to diagnose potential breakdowns in the integrity of randomized trials.[31] The advent of proximal causal inference—a framework for identifying and estimating causal effects under unmeasured confounding by leveraging NCOs and NCEs—has spurred continued interest in negative-control-based methods for causal inference.[32–38]

To illustrate how negative controls may be useful in vaccine studies, we return to the example of influenza vaccination in elderly populations. To investigate whether the association between influenza vaccination and pneumonia/influenza hospitalization and all-cause mortality was contaminated by confounding, Jackson et al.[39] examined the association between influenza vaccination and two negative control outcomes: pneumonia/mortality before the flu season and hospitalization due to trauma. The authors hypothesized that the only way the vaccine could prevent pneumonia, hospitalization, or mortality was by reducing likelihood of influenza-related illnesses. Therefore, the influenza vaccine should have no causal effect on pneumonia/mortality before influenza started circulating or on completely unrelated outcomes like traumatic injuries. Nevertheless, the persistent association between influenza vaccination and both NCOs suggested that residual unmeasured confounding accounted for a substantial portion of the observed association.

Negative controls can also help control for bias in the design phase of real-world vaccine effectiveness studies. The test negative design (TND) is a variant of a case control design which enrolls individuals who seek testing in response to symptoms consistent with the primary infection outcome. Vaccine effectiveness is then estimated by comparing vaccination statuses between individuals with positive test results for the target pathogen (cases) and those with negative test results (noncases). Hence, TNDs are very resource efficient studies that can be readily embedded in healthcare systems. To obtain reliable estimates of vaccine effectiveness, most TND analyses require the assumption that vaccination does not causally affect the noncases,[40–45] which essentially assumes that the noncases are valid negative control outcomes. Recent work has justified VE estimates from TNDs under the assumption that noncases are valid negative controls that capture the magnitude of confounding on the primary endpoint.[46]

We argue that vaccine studies are a particularly compelling setting to study negative controls for several reasons, which we outline below.

First, vaccines are often designed to elicit protection against a specific target pathogen. Vaccines prevent infection/disease by inducing effector mechanisms (i.e., antibodies, memory B-cells, and CD8/CD4 T-cell responses) capable of rapidly controlling replicating pathogens or inactivating their toxic components.[47] These effector functions are highly specific to the target pathogen and, in some cases, it is highly plausible that vaccination will not impact pathogens that are antigenically distinct from the vaccine's target. This is supported by secondary analyses of COVID-19[45,48] and

HPV[49,50] vaccine efficacy studies which found no vaccine effects on off-target infection endpoints in randomized trials. Some possible exceptions include BCG, measles-containing, smallpox, and oral poliovirus vaccines, for which some evidence exists supporting off-target effects.[47] However, for vaccines with highly antigen-specific targets, off-target biological effects are unlikely and the negative control assumption is supported by biological knowledge.

Second, compelling negative controls can be readily conceptualized and collected, sometimes with low additional effort, in vaccine studies. Consider the example of a hypothetical COVID-19 vaccine effectiveness study. Many determinants of COVID-19 vaccination such as healthcare access, occupation, patient and provider beliefs, and perceived risk are impossible to measure directly but are expected to influence uptake of other vaccines (e.g., influenza). These "irrelevant" vaccines represent natural candidates for negative control exposures (NCEs), as they are shaped by the same confounding structure but are not expected to causally affect COVID-19 outcomes. Ascertaining the irrelevant vaccination status of participants can be accomplished with low additional effort. Analogously, determinants of COVID-19 disease risk such as exposure patterns, immune function, and non-pharmaceutical intervention adherence are difficult to measure directly, but are expected to influence the risk of other community-transmitted respiratory infections. Such outcomes represent compelling negative control outcomes (NCOs), as they share common causes with the primary endpoint but should not be affected by the COVID-19 vaccine. These "irrelevant" infection outcomes can often be ascertained through routine monitoring, since participants presenting with respiratory symptoms can be tested for a broader panel of other pathogens.

Finally, the abundance of randomized vaccine trials offer a unique opportunity to empirically test negative control assumptions.[45,48,49] Because vaccination is randomized, any observed association between vaccination and a candidate negative control outcome (or between a candidate negative control exposure and primary outcome) suggests that the candidate negative control is affected by vaccination and is therefore not a valid negative control. In contrast, a small or null association can serve as justification for the negative control assumption. Such tests can help "validate" candidate negative controls and inform their use in secondary analyses and future observational studies.

Motivated by the promise of leveraging negative controls to address methodological challenges in vaccine evaluation, this dissertation develops new statistical methods for improving vaccine effectiveness estimation in randomized trials and observational studies. In Part 1 (Chapters 2 and

3), we will discuss the challenge of estimating how vaccine effectiveness wanes over time and briefly introduce how negative controls infections can help detect and correct for biases in these estimates. In Part 2 (Chapters 4 and 5), we will discuss how negative controls can be used to enhance the efficiency of early and late-phase randomized vaccine trials using methods for covariate adjustment.

1.2 Part 1: Reliable evaluation of time-varying vaccine effectiveness

Summary: Chapters 2 and 3 of my dissertation will detail two projects aiming to improve evaluation of how vaccine effectiveness wanes over time.

A key goal of vaccine research is to develop novel vaccines or adapt existing vaccines to confer long-lasting protection against infection and disease. Understanding how vaccine effectiveness (VE) changes over time is essential for guiding vaccine development and public health decision making, such as when to offer booster vaccinations to the public. For example, evidence of waning influenza vaccine effectiveness supports the practice of vaccinating high-risk groups close to the start of flu season.[51,52] Understanding time-varying VE, or vaccine effectiveness in preventing infection or disease X-days after vaccination, is of interest to both vaccine science and public health.

We take as a given that any meaningful notion of vaccine effectiveness should compare vaccinated and unvaccinated individuals who are otherwise exchangeable, informally meaning the distribution of risk factors are balanced between the groups. Many practitioners estimate time-varying VE by fitting a Cox proportional hazards model with a coefficient that varies in time-since-vaccination.[53–58] While popular among practitioners, time-varying VE estimates based on hazard functions are vulnerable to bias resulting from both structural features of the hazard function and aspects of study design.

Even in randomized placebo-controlled trials, time-varying hazard ratios can provide misleading evidence of VE waning due to a phenomenon referred to as "depletion of susceptibles bias" among epidemiologists.[59,60] In Figure 1.1, we illustrate how this phenomenon arises in a randomized vaccine trial conducted in a population with latent infection risk heterogeneity. For simplicity, we consider of a study population consisting of "low-risk" (blue) and "high-risk" (orange) individuals. In reality, infection risk varies from person-to-person on a continuum, but we consider this simple risk model to clarify how bias arises. Randomization ensures exchangeability—that is, a balance of low-and high-risk individuals between treatment arms—at baseline. To assess VE waning, a common approach is

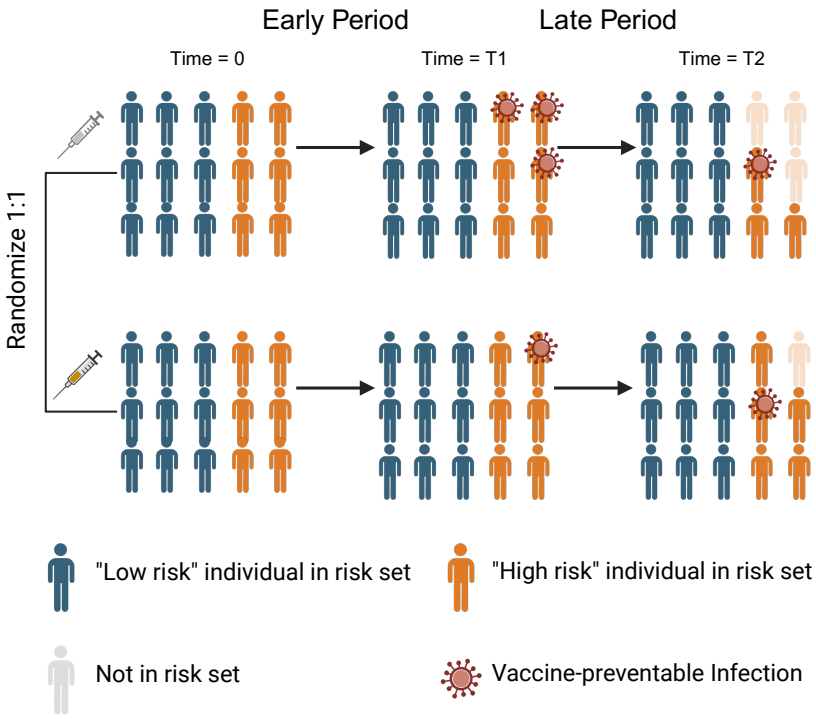

Figure 1.1: Schematic illustrating how depletion of susceptibles bias arises in a randomized vaccine trial.

to compare the hazard ratio during an early period (0 to T1) to the hazard ratio during a late period (T1 to T2). By definition, the late period hazard ratio compares infection rates among vaccinated and unvaccinated persons. If the vaccine is protective, high-risk unvaccinated participants will become infected more rapidly during the early period than high-risk vaccinated participants, leading to imbalance in average infection risk between vaccinated and unvaccinated cohorts who avoided early-period infection. In Figure 1.1, the vaccinated cohort who avoided early period infection had higher average infection risk (i.e., $5/14 \approx 36\%$ orange) than the analogous unvaccinated cohort (i.e., $3/12 = 25\%$ orange). As a result, time-varying hazard ratios can become distorted by time-varying imbalances in underlying infection risk, which leads to exaggerated magnitude of waning VE over time. Using a potential outcomes framework, Hernán[61] clarified that depletion of susceptibles bias is a structural limitation of the hazard function itself, which he termed a "built-in selection bias." In an effort to mitigate depletion of susceptibles, analysts typically adjust for known risk factors, for

example by using serological testing to account for natural immunity from prior infections.[62] Others have proposed alternative vaccine trial designs[59,60] and causal estimands[63] that avoid differential depletion of susceptibles by exploiting natural or hypothetical interventions that delay exposure to the target pathogen until a future date. While these approaches are promising, they require strong assumptions that are unlikely to be satisfied in many contexts.

In some circumstances, conducting long, randomized vaccine trials to evaluate waning VE may not be feasible due to resource limitations, lack of clinical equipoise, or pathogen evolution in the population.[64] As an example, we appeal to COVID-19 vaccine efficacy trials to underscore these constraints. During the COVID-19 pandemic, randomized trials of COVID-19 vaccines had an ethical imperative to unblind participants and offer them vaccination once efficacy has been demonstrated, which complicated the task of evaluating long-term efficacy.[53,54] The availability of COVID-19 vaccines off study could contaminate and complicate the task of estimating vaccine efficacy.[65] Ongoing evolution of the SARS-CoV-2 virus can render VE estimates from historical trials outdated. Conducting a new randomized trial every time a new viral strain emerges would be extremely time-consuming, resource-intensive, and unsuitable for guiding public health responses.

Compared to randomized trials, observational studies are more resource-efficient and offer rich data on how vaccine effectiveness varies across different settings, variants, and over time. However, VE estimates from observational studies are particularly vulnerable to bias. Myriad sources of confounding can impact observational vaccine studies[66,67] including differences in behaviors (such as mask-wearing and social distancing) influencing contact rates with infected community members,[54,68] immune function influencing per-contact susceptibility to infection,[62,68] and healthcare seeking behaviors impacting case ascertainment.[39,69] These concerns are not theoretical: several observational studies in winter 2021 reported low and negative VE estimates, which are hypothesized to be the result of hidden bias.[20]

Common strategies for bias mitigation like observed confounder adjustment are ill-suited to vaccine studies because the sources of bias are hard to measure, overlapping, and potentially time-varying. For example, a potential important confounder is an individual's level of "exposure", or their rate of contact with infected individuals in the community at a particular time.[68] While exposure can be conceptualized and qualitatively described, accurately measuring and incorporating it into statistical models is difficult in practice. Exposure is a multifactorial, time-varying process

that is ill-defined for community-transmitted respiratory pathogens where the only requirement for transmission is spatiotemporal proximity.[70] In practice, analysts typically measure and adjust for demographic factors (e.g., age, employment as a healthcare worker, etc.) with the hope that they are correlated with true unmeasurable confounders.

In Chapters 2 and 3 of this dissertation, we will explore the potential of using negative control infection outcomes to improve time-varying VE estimation by addressing bias arising from both depletion of susceptibles and confounding. In Chapter 2, we focus on the task of bias detection in time-varying VE estimates, which can can be used to identify instances where naïve Cox regression models produce misleading evidence of waning vaccine effects. In Chapter 3, we consider how to identify and estimate a hazard-based, time-varying VE parameter while accounting for unmeasured sources of bias. Both chapters will appeal to the Coronavirus Vaccine Efficacy (COVE; NCT04470427) trial of the mRNA-1273 COVID-19 vaccine as a case study to motivate and illustrate our methods.[3,4]

1.3 Part 2: Improving the efficiency of randomized vaccine trials

Executive Summary: Chapters 4 and 5 of my dissertation will detail two projects that aim to improve the efficiency of randomized trials evaluating vaccines or other immunogens by adjusting for negative control outcomes.

It is well understood that adjusting for baseline covariates can improve efficiency and power in randomized trials by reducing the uncertainty in treatment effect estimates and accounting for finite-sample bias caused by chance covariate imbalances.[71–76] Previous work has proposed covariate-adjusted treatment effect estimators that have theoretically appealing properties, including guarantees of no efficiency loss in large samples and valid inference under a wide range of randomization schemes. Recent work has extended results for continuous outcomes to new estimands tailored to binary,[77], ordinal,[78], survival,[79,80] and recurrent event outcomes.[81] Many of these methods also accommodate the use of nonlinear regression and machine learning models to aid efficiency of treatment effect estimators.[82,83] These approaches exhibit promise to enhance efficiency of randomized trials for biomedical interventions.

From a practical perspective, a key consideration is which covariate(s) to adjust for when estimating the treatment effect of interest. To achieve the greatest improvement in precision and

power, one should adjust for covariates that are believed to be highly correlated with the primary outcome.[84,85] If covariates are highly prognostic, substantial gains in precision and power can be obtained. Alternatively, if the analyst measures and adjusts for weakly prognostic factors, precision gains will be negligible, and precision loss can even occur in finite samples.[84,85] Regulatory guidance supports adjustment for a small number of covariates (relative to the sample size) that are believed to be strongly associated with the primary outcome justified on the basis of scientific knowledge or data from historical trials.[85,86]

However, in vaccine studies, there is often uncertainty around which covariates are prognostic for the primary outcome of interest. In Part 2 of this dissertation, we examine the unique challenges facing covariate adjustment at two key stages of vaccine development—early-phase and late-phase studies. We outline these challenges below and examine them in depth in the subsequent chapters.

In most early-phase vaccine studies, the primary objective is to evaluate the vaccine's safety and immunogenicity, or the immune response elicited by vaccination. Early-phase vaccine trials often have limited sample sizes and imbalanced randomization ratios used to maximize data on vaccine safety. The downstream consequences of these constraints is poor precision in estimates of vaccine-induced immune responses and low power to detect differences in immune responses between different vaccine regimens. In principle, adjustment for baseline covariates is an appealing strategy for improving the precision of vaccine-induced immune response estimates in these resource-constrained studies. However, several factors complicate the application of covariate adjustment in early-phase vaccine studies. First, the validity and theoretical benefits of covariate adjustment (e.g., "guaranteed efficiency gain") are based on asymptotic regimes that may not apply in early-phase studies. For example, even if covariate adjustment cannot harm large-sample efficiency, adjusting for weakly prognostic covariates can lead to precision loss in finite-samples.[84,85] Second, prior research has produced a relative paucity of reliable baseline predictors of vaccine-induced immune responses. A few studies have found modest correlations between age, sex at birth, BMI, and baseline immune responses with post-vaccination T-cell and cytokine responses against HIV-1,[87–90] although high-quality predictors of immune responses remain elusive. The lack of readily available baseline immunogenicity predictors has encouraged the use of systems biology approaches to infer immune responses from high-dimensional immunological readouts.[91–96] Hence, the benefits of covariate adjustment in early-phase randomized trials is limited by incomplete knowledge of

prognostic factors and questionable applicability of asymptotic guarantees to finite-sample settings.

In contrast to early-phase studies, the primary objective of most late-phase vaccine trials is to evaluate the vaccine efficacy, or its ability to prevent infection or disease caused by the target pathogen. Late-phase vaccine trials are expensive and resource-intensive to conduct, particularly for pathogens like HIV-1 where infection outcomes are rare and vaccines may confer only partial protection against infection.[97] In such cases, maximizing the probability of the trial's technical success (i.e., power) given the available resources is critical. Late-phase trials provide a natural setting for covariate-adjusted estimators, as their large sample sizes better match the asymptotic regimes under which these methods are justified. However, factors prognostic for infection risk like "exposure" are exceptionally difficult to measure, while easily measured baseline variables (age, sex, etc.) are typically weakly prognostic.[98–102] Therefore, the realized benefit of covariate adjustment in late-phase vaccine trials is limited by incomplete knowledge and/or measurement of the relevant factors that determine infection risk.

Given incomplete knowledge or inability to measure highly prognostic factors, in vaccine trials, analysts typically avoid adjustment altogether or adjust for a small number of easily-measured, but weakly prognostic variables. In Part 2 of this dissertation, we examine the potential of enhancing the efficiency of vaccine trials by pairing covariate adjustment techniques with negative control outcomes (NCOs), which are presumed to be useful proxies for unmeasured prognostic factors influencing the primary outcomes. In Chapter 4, we investigate adjustment for negative control immune responses (i.e., immune responses against vaccine-mismatched antigens and irrelevant vaccines) as proxies for underlying immune function to improve the precision of vaccine-elicited immune response estimates in early-phase vaccine trials.[103,104] This work emphasizes inference in finite-sample, resource-constrained settings and concludes with reanalyses of two early-phase HIV-1 vaccine studies conducted among HIV-exposed uninfected infants. In Chapter 5, we propose leveraging negative control infection event times to improve the precision of vaccine efficacy estimates and enhance the efficiency of late-phase vaccine trials. In this work, we clarify the assumptions that justify adjustment for negative control infections and address the unique inferential challenges arising from outcome and predictor right-censoring.

Bibliography

[1] World Health Organization. Vaccines and immunization, 2022. URL `https://www.who.int/health-topics/vaccines-and-immunization#tab=tab_1`. Accessed: 2026-04-23.

[2] Andrew J Shattock, Helen C Johnson, So Yoon Sim, Austin Carter, Philipp Lambach, Raymond C W Hutubessy, Kimberly M Thompson, Kamran Badizadegan, Brian Lambert, Matthew J Ferrari, Mark Jit, Han Fu, Sheetal P Silal, Rachel A Hounsell, Richard G White, Jonathan F Mosser, Katy A M Gaythorpe, Caroline L Trotter, Ann Lindstrand, Katherine L O'Brien, and Naor Bar-Zeev. Contribution of vaccination to improved survival and health: modelling 50 years of the expanded programme on immunization. *Lancet*, 403(10441):2307–2316, May 2024.

[3] Lindsey R Baden, Hana M El Sahly, Brandon Essink, Karen Kotloff, Sharon Frey, Rick Novak, David Diemert, Stephen A Spector, Nadine Rouphael, C Buddy Creech, John McGettigan, Shishir Khetan, Nathan Segall, Joel Solis, Adam Brosz, Carlos Fierro, Howard Schwartz, Kathleen Neuzil, Larry Corey, Peter Gilbert, Holly Janes, Dean Follmann, Mary Marovich, John Mascola, Laura Polakowski, Julie Ledgerwood, Barney S Graham, Hamilton Bennett, Rolando Pajon, Conor Knightly, Brett Leav, Weiping Deng, Honghong Zhou, Shu Han, Melanie Ivarsson, Jacqueline Miller, Tal Zaks, and COVE Study Group. Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine. *N. Engl. J. Med.*, 384(5):403–416, February 2021.

[4] Hana M El Sahly, Lindsey R Baden, Brandon Essink, Susanne Doblecki-Lewis, Judith M Martin, Evan J Anderson, Thomas B Campbell, Jesse Clark, Lisa A Jackson, Carl J Fichtenbaum, Marcus Zervos, Bruce Rankin, Frank Eder, Gregory Feldman, Christina Kennelly, Laurie Han-Conrad, Michael Levin, Kathleen M Neuzil, Lawrence Corey, Peter Gilbert, Holly Janes, Dean Follmann, Mary Marovich, Laura Polakowski, John R Mascola, Julie E Ledgerwood, Barney S Graham, Allison August, Heather Clouting, Weiping Deng, Shu Han, Brett Leav, Deb Manzo, Rolando Pajon, Florian Schödel, Joanne E Tomassini, Honghong Zhou, Jacqueline Miller, and COVE Study Group. Efficacy of the mRNA-1273 SARS-CoV-2 vaccine at completion of blinded phase. *N. Engl. J. Med.*, 385(19):1774–1785, November 2021.

[5] Paul T Heath, Eva P Galiza, David N Baxter, Marta Boffito, Duncan Browne, Fiona Burns,

David R Chadwick, Rebecca Clark, Catherine Cosgrove, James Galloway, Anna L Goodman, Amardeep Heer, Andrew Higham, Shalini Iyengar, Arham Jamal, Christopher Jeanes, Philip A Kalra, Christina Kyriakidou, Daniel F McAuley, Agnieszka Meyrick, Angela M Minassian, Jane Minton, Patrick Moore, Imrozia Munsoor, Helen Nicholls, Orod Osanlou, Jonathan Packham, Carol H Pretswell, Alberto San Francisco Ramos, Dinesh Saralaya, Ray P Sheridan, Richard Smith, Roy L Soiza, Pauline A Swift, Emma C Thomson, Jeremy Turner, Marianne E Viljoen, Gary Albert, Iksung Cho, Filip Dubovsky, Greg Glenn, Joy Rivers, Andreana Robertson, Kathy Smith, Seth Toback, and 2019nCoV-302 Study Group. Safety and efficacy of NVX-CoV2373 covid-19 vaccine. *N. Engl. J. Med.*, 385(13):1172–1183, September 2021.

[6] Fernando P Polack, Stephen J Thomas, Nicholas Kitchin, Judith Absalon, Alejandra Gurtman, Stephen Lockhart, John L Perez, Gonzalo Pérez Marc, Edson D Moreira, Cristiano Zerbini, Ruth Bailey, Kena A Swanson, Satrajit Roychoudhury, Kenneth Koury, Ping Li, Warren V Kalina, David Cooper, Robert W Frenck, Jr, Laura L Hammitt, Özlem Türeci, Haylene Nell, Axel Schaefer, Serhat Ünal, Dina B Tresnan, Susan Mather, Philip R Dormitzer, Uğur Şahin, Kathrin U Jansen, William C Gruber, and C4591001 Clinical Trial Group. Safety and efficacy of the BNT162b2 mRNA covid-19 vaccine. *N. Engl. J. Med.*, 383(27):2603–2615, December 2020.

[7] Merryn Voysey, Sue Ann Costa Clemens, Shabir A Madhi, Lily Y Weckx, Pedro M Folegatti, Parvinder K Aley, Brian Angus, Vicky L Baillie, Shaun L Barnabas, Qasim E Bhorat, Sagida Bibi, Carmen Briner, Paola Cicconi, Andrea M Collins, Rachel Colin-Jones, Clare L Cutland, Thomas C Darton, Keertan Dheda, Christopher J A Duncan, Katherine R W Emary, Katie J Ewer, Lee Fairlie, Saul N Faust, Shuo Feng, Daniela M Ferreira, Adam Finn, Anna L Goodman, Catherine M Green, Christopher A Green, Paul T Heath, Catherine Hill, Helen Hill, Ian Hirsch, Susanne H C Hodgson, Alane Izu, Susan Jackson, Daniel Jenkin, Carina C D Joe, Simon Kerridge, Anthonet Koen, Gaurav Kwatra, Rajeka Lazarus, Alison M Lawrie, Alice Lelliott, Vincenzo Libri, Patrick J Lillie, Raburn Mallory, Ana V A Mendes, Eveline P Milan, Angela M Minassian, Alastair McGregor, Hazel Morrison, Yama F Mujadidi, Anusha Nana, Peter J O'Reilly, Sherman D Padayachee, Ana Pittella, Emma Plested, Katrina M

Pollock, Maheshi N Ramasamy, Sarah Rhead, Alexandre V Schwarzbold, Nisha Singh, Andrew Smith, Rinn Song, Matthew D Snape, Eduardo Sprinz, Rebecca K Sutherland, Richard Tarrant, Emma C Thomson, M Estée Török, Mark Toshner, David P J Turner, Johan Vekemans, Tonya L Villafana, Marion E E Watson, Christopher J Williams, Alexander D Douglas, Adrian V S Hill, Teresa Lambe, Sarah C Gilbert, Andrew J Pollard, and Oxford COVID Vaccine Trial Group. Safety and efficacy of the ChAdOx1 nCoV-19 vaccine (AZD1222) against SARS-CoV-2: an interim analysis of four randomised controlled trials in brazil, south africa, and the UK. *Lancet*, 397(10269):99–111, January 2021.

[8] Ann R Falsey, Magdalena E Sobieszczyk, Ian Hirsch, Stephanie Sproule, Merlin L Robb, Lawrence Corey, Kathleen M Neuzil, William Hahn, Julie Hunt, Mark J Mulligan, Charlene McEvoy, Edwin DeJesus, Michael Hassman, Susan J Little, Barbara A Pahud, Anna Durbin, Paul Pickrell, Eric S Daar, Larry Bush, Joel Solis, Quito Osuna Carr, Temitope Oyedele, Susan Buchbinder, Jessica Cowden, Sergio L Vargas, Alfredo Guerreros Benavides, Robert Call, Michael C Keefer, Beth D Kirkpatrick, John Pullman, Tina Tong, Margaret Brewinski Isaacs, David Benkeser, Holly E Janes, Martha C Nason, Justin A Green, Elizabeth J Kelly, Jill Maaske, Nancy Mueller, Kathryn Shoemaker, Therese Takas, Richard P Marshall, Menelas N Pangalos, Tonya Villafana, Antonio Gonzalez-Lopez, and AstraZeneca AZD1222 Clinical Study Group. Phase 3 safety and efficacy of AZD1222 (ChAdOx1 nCoV-19) covid-19 vaccine. *N. Engl. J. Med.*, 385(25):2348–2360, December 2021.

[9] John P A Ioannidis, Angelo Maria Pezzullo, Antonio Cristiano, and Stefania Boccia. Global estimates of lives and life-years saved by COVID-19 vaccination during 2020-2024. *JAMA Health Forum*, 6(7):e252223, July 2025.

[10] Marc Lipsitch and Nir Eyal. Improving vaccine trials in infectious disease emergencies. *Science*, 357(6347):153–156, July 2017.

[11] M E Halloran, M Haber, I M Longini, Jr, and C J Struchiner. Direct and indirect effects in vaccine efficacy and effectiveness. *Am. J. Epidemiol.*, 133(4):323–331, February 1991.

[12] M E Halloran, C J Struchiner, and I M Longini, Jr. Study designs for evaluating different

efficacy and effectiveness aspects of vaccines. *Am. J. Epidemiol.*, 146(10):789–803, November 1997.

[13] S J Day and D G Altman. Statistics notes: blinding in clinical trials and other studies. *BMJ*, 321(7259):504, 2000.

[14] M E Halloran, M Haber, I M Longini, Jr, and C J Struchiner. Direct and indirect effects in vaccine efficacy and effectiveness. *Am. J. Epidemiol.*, 133(4):323–331, February 1991.

[15] Stephen J W Evans and Nicholas P Jewell. Vaccine effectiveness studies in the field. *N. Engl. J. Med.*, 385(7):650–651, August 2021.

[16] Marco Carone, Charles J Wolock, Antonio Olivas-Martinez, Andrea Rotnitzky, and Peter B Gilbert. Immune correlates and vaccine immunobridging: Statistical innovations, challenges, and opportunities. *J. Infect. Dis.*, 232(5):1024–1028, November 2025.

[17] Helen Y Chu, Holly Janes, Marco Carone, Peter B Gilbert, and Stanley Plotkin. Improving the evidence base for COVID-19 vaccines. *Nat. Med.*, 31(11):3595–3596, November 2025.

[18] Vittorio Demicheli, Tom Jefferson, Carlo Di Pietrantonj, Eliana Ferroni, Sarah Thorning, Roger E Thomas, and Alessandro Rivetti. Vaccines for preventing influenza in the elderly. *Cochrane Database Syst. Rev.*, 2(10):CD004876, February 2018.

[19] Lone Simonsen, Robert J Taylor, Cecile Viboud, Mark A Miller, and Lisa A Jackson. Mortality benefits of influenza vaccination in elderly people: an ongoing controversy. *Lancet Infect. Dis.*, 7(10):658–666, October 2007.

[20] Korryn Bodner, Michael A. Irvine, Jeffrey C. Kwong, and Sharmistha Mishra. Observed negative vaccine effectiveness could be the canary in the coal mine for biases in observational COVID-19 studies. *International Journal of Infectious Diseases*, 131:111–114, June 2023. ISSN 12019712. doi: 10.1016/j.ijid.2023.03.022.

[21] Paul R Rosenbaum. The role of known effects in observational studies. *Biometrics*, 45(2): 557–569, 1984.

[22] Paul R Rosenbaum. Detecting bias with confidence in observational studies. *Biometrika*, 79 (2):367–374, 1992.

[23] Marc Lipsitch, Eric Tchetgen Tchetgen, and Ted Cohen. Negative controls: a tool for detecting confounding and bias in observational studies. *Epidemiology*, 21(3):383–388, May 2010.

[24] Stacie B Dusetzina, M Alan Brookhart, and Matthew L Maciejewski. Control outcomes and exposures for improving internal validity of nonrandomized studies. *Health Serv. Res.*, 50(5): 1432–1451, October 2015.

[25] Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. A selective review of negative control methods in epidemiology. *Curr. Epidemiol. Rep.*, 7(4):190–202, December 2020.

[26] Eric Tchetgen Tchetgen. The control outcome calibration approach for causal inference with unobserved confounding. *Am. J. Epidemiol.*, 179(5):633–640, March 2014.

[27] David B Richardson, Dominique Laurier, Mary K Schubauer-Berigan, Eric Tchetgen Tchetgen, and Stephen R Cole. Assessment and indirect adjustment for confounding by smoking in cohort studies using relative hazards models. *Am. J. Epidemiol.*, 180(9):933–940, November 2014.

[28] David B Richardson, Alexander P Keil, Eric Tchetgen Tchetgen, and Glinda Cooper. Negative control outcomes and the analysis of standardized mortality ratios. *Epidemiology*, 26(5):727–732, September 2015.

[29] Tamar Sofer, David B Richardson, Elena Colicino, Joel Schwartz, and Eric J Tchetgen Tchetgen. On negative outcome control of unobserved confounding as a generalization of difference-in-differences. *Stat. Sci.*, 31(3):348–361, September 2016.

[30] Benjamin F. Arnold, Ayse Ercumen, Jade Benjamin-Chung, and John M. Colford. Brief Report: Negative Controls to Detect Selection Bias and Measurement Bias in Epidemiologic Studies. *Epidemiology*, 27(5):637–641, September 2016. ISSN 1044-3983. doi: 10.1097/EDE.0000000000000504. URL `http://journals.lww.com/00001648-201609000-00006`.

[31] Benjamin F Arnold and Ayse Ercumen. Negative control outcomes: A tool to detect bias in randomized trials. *JAMA*, 316(24):2597–2598, December 2016.

[32] Wang Miao, Zhi Geng, and Eric Tchetgen Tchetgen. Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*, 105(4):987–993, December 2018.

[33] Xu Shi, Wang Miao, Jennifer C Nelson, and Eric J Tchetgen Tchetgen. Multiply robust causal inference with double-negative control adjustment for categorical unmeasured confounding. *J. R. Stat. Soc. Series B Stat. Methodol.*, 82(2):521–540, April 2020.

[34] Andrew Ying, Wang Miao, Xu Shi, and Eric J Tchetgen Tchetgen. Proximal causal inference for complex longitudinal studies. *J. R. Stat. Soc. Series B Stat. Methodol.*, March 2023.

[35] Paul N Zivich, Stephen R Cole, Jessie K Edwards, Grace E Mulholland, Bonnie E Shook-Sa, and Eric J Tchetgen Tchetgen. Introducing proximal causal inference for epidemiologists. *Am. J. Epidemiol.*, 192(7):1224–1227, July 2023.

[36] Jiewen Liu, Chan Park, Kendrick Li, and Eric J Tchetgen Tchetgen. Regression-based proximal causal inference. *Am. J. Epidemiol.*, 194(7):kwae370, September 2024.

[37] Eric J Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. An introduction to proximal causal inference. *Stat. Sci.*, 39(3), August 2024.

[38] Yifan Cui, Hongming Pu, Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. Semiparametric proximal causal inference. *J. Am. Stat. Assoc.*, 119(546):1348–1359, April 2024.

[39] Lisa A Jackson, Michael L Jackson, Jennifer C Nelson, Kathleen M Neuzil, and Noel S Weiss. Evidence of bias in estimates of influenza vaccine effectiveness in seniors. *International Journal of Epidemiology*, 35:337–344, 2005.

[40] Michael L Jackson and Jennifer C Nelson. The test-negative design for estimating influenza vaccine effectiveness. *Vaccine*, 31(17):2165–2168, April 2013.

[41] Claire V. Broome, Richard R. Facklam, and David W. Fraser. Pneumococcal Disease after Pneumococcal Vaccination: An Alternative Method to Estimate the Efficacy of Pneumococcal

Vaccine. *New England Journal of Medicine*, 303(10):549–552, September 1980. ISSN 0028-4793, 1533-4406. doi: 10.1056/NEJM198009043031003. URL `http://www.nejm.org/doi/abs/10.1056/NEJM198009043031003`.

[42] G De Serres, D M Skowronski, X W Wu, and C S Ambrose. The test-negative design: validity, accuracy and precision of vaccine efficacy estimates compared to the gold standard of randomised placebo-controlled clinical trials. *Euro Surveill.*, 18(37):20585, September 2013.

[43] Mireille E Schnitzer. Estimands and estimation of COVID-19 vaccine effectiveness under the test-negative design: Connections to causal inference. *Epidemiology*, 33(3):325–333, May 2022.

[44] Cong Jiang, Denis Talbot, Sara Carazo, and Mireille E Schnitzer. A double machine learning approach for the evaluation of COVID-19 vaccine effectiveness under the test-negative design: Analysis of québec administrative data. *Stat. Med.*, 44(5):e70025, February 2025.

[45] Leah I B Andrews, M Elizabeth Halloran, Kathleen M Neuzil, Lars van der Laan, Yunda Huang, Jessica Andriesen, Mayur Patel, Leigh H Fisher, Holly Janes, Nadine Rouphael, Stephen R Walsh, Deborah A Theodore, Hong-Van Tieu, Magdalena Sobieszczyk, Hana M El Sahly, Lindsey R Baden, Ann R Falsey, Thomas B Campbell, Colleen F Kelley, Catherine Mary Healy, Lilly Immergluck, Benjamin Luft, Ian Hirsch, Guy de Bruyn, Carla Truyers, Frances Priddy, Kelsey M Sumner, Brendan Flannery, Dean Follmann, Peter B Gilbert, and COVID-19 Prevention Network (CoVPN). Evaluating the test-negative design for COVID-19 vaccine effectiveness using randomized trial data: A secondary cross-protocol analysis of 5 randomized clinical trials. *JAMA Network Open*, 8(5):e2512763, May 2025.

[46] Christopher B Boyer, Kendrick Qijun Li, Xu Shi, and Eric J Tchetgen Tchetgen. Identification and estimation of vaccine effectiveness in the test-negative design under equi-confounding. *Epidemiology*, 37(1):77–87, January 2026.

[47] Walter A Orenstein, Paul A Offit, Kathryn M Edwards, and Stanley A Plotkin. *Plotkin's vaccines*. Elsevier - Health Sciences Division, Philadelphia, PA, 8 edition, May 2023.

[48] Ethan Ashby, Holly Janes, Dean Follmann, Peter B Gilbert, Honghong Zhou, Xiaowei Wang, Bethany Girard, Frances Priddy, James G Kublin, Lawrence Corey, Kathleen M Neuzil, Lindsey R Baden, Hana M El Sahly, Bo Zhang, and COVE study group. Validating and leveraging non-SARS-CoV-2 respiratory infection as a negative control outcome in a phase 3 COVID-19 vaccine trial with extended observational follow-up. *American Journal of Epidemiology*, 195 (1):168–177, January 2026.

[49] Joseph E Tota, Frank Struyf, Joshua N Sampson, Paula Gonzalez, Martin Ryser, Rolando Herrero, John Schussler, Naveen Karkada, Ana Cecilia Rodriguez, Nicolas Folschweiller, Carolina Porras, Mark Schiffman, John T Schiller, Wim Quint, Aimée R Kreimer, Cosette M Wheeler, Allan Hildesheim, and Costa Rica Vaccine Trial and PATRICIA Study. Efficacy of the AS04-adjuvanted HPV16/18 vaccine: Pooled analysis of the costa rica vaccine and PATRICIA randomized controlled trials. *J. Natl. Cancer Inst.*, 112(8):818–828, August 2020.

[50] Lola Etievant, Joshua N. Sampson, and Mitchell H. Gail. Increasing efficiency and reducing bias when assessing HPV vaccination efficacy by using nontargeted HPV strains. *Biometrics*, 79:1534–1545, 2023. doi: 10.1111/biom.13663. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/biom.13663`.

[51] Edward A. Belongia, Maria E. Sundaram, David L. McClure, Jennifer K. Meece, Jill Ferdinands, and Jeffrey J VanWormer. Waning vaccine protection against influenza a (h3n2) illness in children and older adults during a single season. *Vaccine*, 33(1):246–251, 2015. ISSN 0264-410X. doi: https://doi.org/10.1016/j.vaccine.2014.06.052. URL `https://www.sciencedirect.com/science/article/pii/S0264410X14008536`.

[52] Jill M. Ferdinands, Alicia M. Fry, Sue Reynolds, Joshua G. Petrie, Brendan Flannery, Michael L Jackson, and Edward A Belongia. Intraseason waning of influenza vaccine protection: Evidence from the us influenza vaccine effectiveness network, 2011–2012 through 2014–2015. *Clinical Infectious Diseases*, 64(5):544–550, 02 2017. ISSN 1058-4838. doi: 10.1093/cid/ciw816. URL `https://doi.org/10.1093/cid/ciw816`.

[53] Jonathan Fintzi and Dean Follmann. Assessing vaccine durability in randomized trials following placebo crossover. *Stat. Med.*, 40(27):5983–6007, November 2021.

[54] Anastasios A Tsiatis and Marie Davidian. Estimating vaccine efficacy over time after a randomized study is unblinded. *Biometrics*, 78(3):825–838, September 2022.

[55] Dan-Yu Lin, Donglin Zeng, and Peter B Gilbert. Evaluating the long-term efficacy of coronavirus disease 2019 (COVID-19) vaccines. *Clin. Infect. Dis.*, 73(10):1927–1939, November 2021.

[56] Dan-Yu Lin, Lindsey R Baden, Hana M El Sahly, Brandon Essink, Kathleen M Neuzil, Lawrence Corey, Jacqueline Miller, and COVE Study Group. Durability of protection against symptomatic COVID-19 among participants of the mRNA-1273 SARS-CoV-2 vaccine trial. *JAMA Netw. Open*, 5(6):e2215984, June 2022.

[57] Dan-Yu Lin, Yu Gu, Bradford Wheeler, Hayley Young, Shannon Holloway, Shadia-Khan Sunny, Zack Moore, and Donglin Zeng. Effectiveness of covid-19 vaccines over a 9-month period in north carolina. *N. Engl. J. Med.*, 386(10):933–941, March 2022.

[58] Dan-Yu Lin, Yi Du, Yangjianchen Xu, Sai Paritala, Matthew Donahue, and Patrick Maloney. Durability of XBB.1.5 vaccines against omicron subvariants. *N. Engl. J. Med.*, 390(22):2124–2127, June 2024.

[59] Marc Lipsitch. Challenges of vaccine effectiveness and waning studies. *Clin. Infect. Dis.*, 68 (10):1631–1633, May 2019.

[60] G Thomas Ray, Ned Lewis, Nicola P Klein, Matthew F Daley, Marc Lipsitch, and Bruce Fireman. Depletion-of-susceptibles bias in analyses of intra-season waning of influenza vaccine effectiveness. *Clin. Infect. Dis.*, 70(7):1484–1486, March 2020.

[61] Miguel A. Hernán. The hazards of hazard ratios. *Epidemiology*, 21(1):13–15, January 2010.

[62] Rebecca Kahn, Stephanie J Schrag, Jennifer R Verani, and Marc Lipsitch. Identifying and alleviating bias due to differential depletion of susceptible people in postmarketing evaluations of COVID-19 vaccines. *Am. J. Epidemiol.*, 191(5):800–811, March 2022.

[63] Matias Janvin and Mats J Stensrud. Quantification of vaccine waning as a challenge effect. *J. Am. Stat. Assoc.*, pages 1–11, December 2024.

[64] Mona N. Kanaan and C. Paddy Farrington. Estimation of Waning Vaccine Efficacy. *Journal of the American Statistical Association*, 97(458):389–397, 2002. ISSN 0162-1459. URL `https://www.jstor.org/stable/3085652`. Publisher: [American Statistical Association, Taylor & Francis, Ltd.].

[65] Kathryn E Stephenson, Jasmine R Marcelin, Audrey E Pettifor, Holly Janes, Elizabeth Brown, Moni Neradilek, Catherine Yen, Jessica Andriesen, Nicole Grunenberg, Nicole Espy, Meg Trahey, Rebecca S B Fischer, Christopher A DeSouza, Joanna L Shisler, Elizabeth Connick, Eric R Houpt, Helen Y Chu, Russel J McCulloh, Sylvia Becker-Dreps, Nadja A Vielot, Corey A Kalbaugh, Kartik Cherabuddi, Karen M Krueger, Molly Rosenberg, Richard N Greenberg, Arnel Joaquin, Lilly Cheng Immergluck, Lawrence Corey, and James G Kublin. Efficacy of messenger RNA-1273 against severe acute respiratory syndrome coronavirus 2 acquisition in young adults from march to december 2021. *Open Forum Infect. Dis.*, 10(11): ofad511, November 2023.

[66] World Health Organization. Evaluation of COVID-19 vaccine effectiveness: Interim Guidance, March 2021.

[67] World Health Organization. Evaluation of COVID-19 vaccine effectiveness in a changing landscape of COVID-19 epidemiology and vaccination: Interim Guidance, October 2022.

[68] M.E. Halloran, I.M. Longini, and C.J. Struchiner. Estimability and interpretation of vaccine efficacy using frailty mixing models. *American Journal of Epidemiology*, 144:83–97, 1996.

[69] Michael L Jackson and Jennifer C Nelson. The test-negative design for estimating influenza vaccine effectiveness. *Vaccine*, 31(17):2165–2168, 2013.

[70] Steffen Unkel, C. Paddy Farrington, Heather J, Whitaker, and Richard Peabody. Time varying frailty models and the estimation of heterogeneities in transmission of infectious diseases. *Journal of the Royal Statistical Society: Series C*, 63(1):141–158, 2014.

[71] Li Yang and Anastasios A. Tsiatis. Efficiency study of estimators for a treatment effect in a pretest-posttest trial. *The American Statistician*, 55:314–321, 2001.

[72] Anastasios A. Tsiatis, Marie Davidian, Min Zhang, and Xiaomin Lu. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Statistics in medicine*, 27(23):4658–4677, October 2008. ISSN 0277-6715. doi: 10.1002/sim.3113. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2562926/`.

[73] Min Zhang, Anastasios A. Tsiatis, and Marie Davidian. Improving Efficiency of Inferences in Randomized Clinical Trials Using Auxiliary Covariates. *Biometrics*, 64 (3):707–715, 2008. ISSN 1541-0420. doi: 10.1111/j.1541-0420.2007.00976.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1541-0420.2007.00976.x.

[74] Winston Lin. Agnostic notes on regression adjustments to experimental data: Reexamining freedman's critique. *Annals of Applied Statistics*, 7:295–318, 2013.

[75] Brennan C Kahan, Vipul Jairath, Caroline J Doré, and Tim P Morris. The risks and rewards of covariate adjustment in randomized trials: an assessment of 12 outcomes from 8 studies. *Trials*, 15(1):139, April 2014.

[76] Ting Ye, Jun Shao, Yanyao Yi, and Qingquan Zhao. Toward better practice of covariate adjustment in analyzing randomized clinical trials. *Journal of the American Statistical Association*, 2022.

[77] Xiaoyu Qiu, Yuhan Qian, Jaehwan Yi, Jinqiu Wang, Yu Du, Yanyao Yi, and Ting Ye. Clarifying the role of the Mantel-Haenszel risk difference estimator in randomized clinical trials. August 2024.

[78] Zhilan Lou, Jun Shao, Ting Ye, Tuo Wang, Yanyao Yi, and Yu Du. Covariate adjustment for wilcoxon two sample statistic and test. February 2026.

[79] Ting Ye and Jun Shao. Robust tests for treatment effect in survival analysis under covariate-adaptive randomization. *J. R. Stat. Soc. Series B Stat. Methodol.*, 82(5):1301–1323, December 2020.

[80] Ting Ye, Jun Shao, and Yanyao Yi. Covariate-adjusted log-rank test: guaranteed efficiency gain and universal applicability. *Biometrika*, 111(2):691–705, May 2024.

[81] Jiren Sun, Tuo Wang, Yanyao Yi, Ting Ye, Jun Shao, and Yu Du. Improve the precision of area under the curve estimation for recurrent events through covariate adjustment. *Stat. Med.*, 44(15-17):e70187, July 2025.

[82] Marlena S. Bannick, Jun Shao, Jingyi Liu, Yu Du, Yanyao Yi, and Ting Ye. A general form of covariate adjustment in randomized clinical trials. *Biometrika*, 112, 2025. doi: https://doi.org/10.1093/biomet/asaf029.

[83] Kevin Guo and Guillaume Basse. The generalized oaxaca-blinder estimator. *J. Am. Stat. Assoc.*, 118(541):524–536, January 2023.

[84] Kelly Van Lancker, Frank Bretz, and Oliver Dukes. Covariate adjustment in randomized controlled trials: General concepts and practical considerations. *Clin. Trials*, 21(4):399–411, August 2024.

[85] U.S. Food and Drug Administration. Adjusting for covariates in randomized clinical trials for drugs and biological products: Draft guidance for industry, May 2021. URL `https://www.fda.gov/media/148910/download`. Draft guidance.

[86] European Medicines Agency. Guideline on adjustment for baseline covariates in clinical trials, March 2015. URL `https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-adjustment-baseline-covariates-clinical-trials_en.pdf`.

[87] Xia Jin, Cecilia Morgan, Xuesong Yu, Stephen DeRosa, Georgia D Tomaras, David C Montefiori, James Kublin, Larry Corey, Michael C Keefer, and NIAID HIV Vaccine Trials Network. Multiple factors affect immunogenicity of DNA plasmid HIV vaccines in human clinical trials. *Vaccine*, 33(20):2347–2353, May 2015.

[88] Susan P Buchbinder, Devan V Mehrotra, Ann Duerr, Daniel W Fitzgerald, Robin Mogg, David Li, Peter B Gilbert, Javier R Lama, Michael Marmor, Carlos Del Rio, M Juliana McElrath, Danilo R Casimiro, Keith M Gottesdiener, Jeffrey A Chodakewitz, Lawrence Corey, Michael N Robertson, and Step Study Protocol Team. Efficacy assessment of a cell-mediated immunity HIV-1 vaccine (the step study): a double-blind, randomised, placebo-controlled, test-of-concept trial. *Lancet*, 372(9653):1881–1893, November 2008.

[89] Adamson Blythe, Tomer Hertz, Ann Duerr, Cecilia Morgan, Michael Keefer, Gavin Churchyard, Barney Graham, P.B. Gilbert, and James Kulbin. Baseline predictors of immunogenicity in hvtn 204. In *Proceedings of the 2013 AIDS Vaccine Conference*, pages 7–10, Barcelona, Spain, 2013.

[90] Yunda Huang, Yuanyuan Zhang, Kelly E Seaton, Stephen De Rosa, Jack Heptinstall, Lindsay N Carpp, April Kaur Randhawa, Lyle R McKinnon, Paul McLaren, Edna Viegas, Glenda E Gray, Gavin Churchyard, Susan P Buchbinder, Srilatha Edupuganti, Linda-Gail Bekker, Michael C Keefer, Mina C Hosseinipour, Paul A Goepfert, Kristen W Cohen, Brian D Williamson, M Juliana McElrath, Georgia D Tomaras, Juilee Thakar, James J Kobie, and NIAID HIV Vaccine Trials Network (HVTN) 073/South African AIDS Vaccine Initiative (SAAVI) 102, HVTN 086/SAAVI 103, HVTN 094, HVTN 097, HVTN 098, HVTN 100, HVTN 105, HVTN 111, and HVTN 205 Protocol Leadership Teams. Baseline host determinants of robust human HIV-1 vaccine-induced immune responses: A meta-analysis of 26 vaccine regimens. *EBioMedicine*, 84(104271):104271, October 2022.

[91] John S Tsang, Pamela L Schwartzberg, Yuri Kotliarov, Angelique Biancotto, Zhi Xie, Ronald N Germain, Ena Wang, Matthew J Olnes, Manikandan Narayanan, Hana Golding, Susan Moir, Howard B Dickler, Shira Perl, Foo Cheung, Baylor HIPC Center, and CHI Consortium. Global analyses of human immune variation reveal baseline predictors of post-vaccination responses. *Cell*, 157(2):499–513, April 2014.

[92] HIPC-CHI Signatures Project Team and HIPC-I Consortium. Multicohort analysis reveals baseline transcriptional predictors of influenza vaccination responses. *Sci. Immunol.*, 2(14): eaal4656, August 2017.

[93] Slim Fourati, Aarthi Talla, Mehrad Mahmoudian, Joshua G Burkhart, Riku Klén, Ricardo Henao, Thomas Yu, Zafer Aydın, Ka Yee Yeung, Mehmet Eren Ahsen, Reem Almugbel, Samad Jahandideh, Xiao Liang, Torbjörn E M Nordling, Motoki Shiga, Ana Stanescu, Robert Vogel, Respiratory Viral DREAM Challenge Consortium, Gaurav Pandey, Christopher Chiu, Micah T McClain, Christopher W Woods, Geoffrey S Ginsburg, Laura L Elo, Ephraim L Tsalik, Lara M Mangravite, and Solveig K Sieberts. A crowdsourced analysis to identify ab

initio molecular signatures predictive of susceptibility to viral infection. *Nat. Commun.*, 9(1): 4418, October 2018.

[94] Slim Fourati, Lewis E Tomalin, Matthew P Mulè, Daniel G Chawla, Bram Gerritsen, Dmitry Rychkov, Evan Henrich, Helen E R Miller, Thomas Hagan, Joann Diray-Arce, Patrick Dunn, Human Immunology Project Consortium (HIPC), Ofer Levy, Raphael Gottardo, Minnie M Sarwal, John S Tsang, Mayte Suárez-Fariñas, Bali Pulendran, Steven H Kleinstein, and Rafick-Pierre Sékaly. Pan-vaccine analysis reveals innate immune endotypes predictive of antibody responses to vaccination. *Nat. Immunol.*, 23(12):1777–1787, December 2022.

[95] Pramod Shinde, Ferran Soldevila, Joaquin Reyna, Minori Aoki, Mikkel Rasmussen, Lisa Willemsen, Mari Kojima, Brendan Ha, Jason A Greenbaum, James A Overton, Hector Guzman-Orozco, Somayeh Nili, Shelby Orfield, Jeremy P Gygi, Ricardo da Silva Antunes, Alessandro Sette, Barry Grant, Lars Rønn Olsen, Anna Konstorum, Leying Guan, Ferhat Ay, Steven H Kleinstein, and Bjoern Peters. A multi-omics systems vaccinology resource to develop and test computational models of immunity. *Cell Rep. Methods*, 4(3):100731, March 2024.

[96] Pramod Shinde, Lisa Willemsen, Michael Anderson, Minori Aoki, Saonli Basu, Julie G Burel, Peng Cheng, Souradipto Ghosh Dastidar, Aidan Dunleavy, Tal Einav, Jamie Forschmiedt, Slim Fourati, Javier Garcia, William Gibson, Jason A Greenbaum, Leying Guan, Weikang Guan, Jeremy P Gygi, Brendan Ha, Joe Hou, Jason Hsiao, Yunda Huang, Rick Jansen, Bhargob Kakoty, Zhiyu Kang, James J Kobie, Mari Kojima, Anna Konstorum, Jiyeun Lee, Sloan A Lewis, Aixin Li, Eric F Lock, Jarjapu Mahita, Marcus Mendes, Hailong Meng, Aidan Neher, Somayeh Nili, Lars Rønn Olsen, Shelby Orfield, James A Overton, Nidhi Pai, Cokie Parker, Brian Qian, Mikkel Rasmussen, Joaquin Reyna, Eve Richardson, Sandra Safo, Josey Sorenson, Aparna Srinivasan, Nicola Thrupp, Rashmi Tippalagama, Raphael Trevizani, Steffen Ventz, Jiuzhou Wang, Cheng-Chang Wu, Ferhat Ay, Barry Grant, Steven H Kleinstein, and Bjoern Peters. Putting computational models of immunity to the test-an invited challenge to predict b.pertussis vaccination responses. *PLoS Comput. Biol.*, 21(3):e1012927, March 2025.

[97] Lawrence Corey, Peter B Gilbert, Michal Juraska, David C Montefiori, Lynn Morris, Shelly T

Karuna, Srilatha Edupuganti, Nyaradzo M Mgodi, Allan C deCamp, Erika Rudnicki, Yunda Huang, Pedro Gonzales, Robinson Cabello, Catherine Orrell, Javier R Lama, Fatima Laher, Erica M Lazarus, Jorge Sanchez, Ian Frank, Juan Hinojosa, Magdalena E Sobieszczyk, Kyle E Marshall, Pamela G Mukwekwerere, Joseph Makhema, Lindsey R Baden, James I Mullins, Carolyn Williamson, John Hural, M Juliana McElrath, Carter Bentley, Simbarashe Takuva, Margarita M Gomez Lorenzo, David N Burns, Nicole Espy, April K Randhawa, Nidhi Kochar, Estelle Piwowar-Manning, Deborah J Donnell, Nirupama Sista, Philip Andrew, James G Kublin, Glenda Gray, Julie E Ledgerwood, John R Mascola, Myron S Cohen, and HVTN 704/HPTN 085 and HVTN 703/HPTN 081 Study Teams. Two randomized trials of neutralizing antibodies to prevent HIV-1 acquisition. *N. Engl. J. Med.*, 384(11):1003–1014, March 2021.

[98] M Elizabeth Halloran, Ira M Longini, and Claudio Jose Struchiner. *Design and analysis of vaccine studies*. Statistics for Biology and Health. Springer, New York, NY, 2010 edition, November 2009.

[99] Merlin L Robb, Supachai Rerks-Ngarm, Sorachai Nitayaphan, Punnee Pitisuttithum, Jaranit Kaewkungwal, Prayura Kunasol, Chirasak Khamboonruang, Prasert Thongcharoen, Patricia Morgan, Michael Benenson, Robert M Paris, Joseph Chiu, Elizabeth Adams, Donald Francis, Sanjay Gurunathan, Jim Tartaglia, Peter Gilbert, Don Stablein, Nelson L Michael, and Jerome H Kim. Risk behaviour and time as covariates for efficacy of the HIV vaccine regimen ALVAC-HIV (vCP1521) and AIDSVAX B/E: a post-hoc analysis of the thai phase 3 efficacy trial RV 144. *Lancet Infect. Dis.*, 12(7):531–537, July 2012.

[100] Peter B Gilbert. Ongoing vaccine and monoclonal antibody HIV prevention efficacy trials and considerations for sequel efficacy trial designs. *Stat. Commun. Infect. Dis.*, 11(1):20190003, July 2019.

[101] Justin J O'Hagan, Marc Lipsitch, and Miguel A Hernán. Estimating the per-exposure effect of infectious disease interventions. *Epidemiology*, 25(1):134–138, January 2014.

[102] Mats J Stensrud and Louisa Smith. Identification of vaccine effects when exposure status is unknown. *Epidemiology*, 34(2):216–224, March 2023.

[103] Dean Follmann. Augmented designs to assess immune response in vaccine trials. *Biometrics*, 62(4):1161–1169, December 2006.

[104] Ying Huang, Shomoita Alam, Erica Andersen-Nissen, Lindsay N Carpp, One B Dintwe, Britta S Flach, Nicole Grunenberg, Fatima Laher, Stephen C De Rosa, Guido Ferrari, Craig Innes, Linda-Gail Bekker, James G Kublin, M Juliana McElrath, Georgia D Tomaras, Glenda E Gray, and Peter B Gilbert. Non-HIV vaccine-induced immune responses as potential baseline immunogenicity predictors of ALVAC-HIV and AIDSVAX B/E-induced immune responses. *Viruses*, 16(9):1365, August 2024.

Chapter 2

VALIDATING AND LEVERAGING NEGATIVE CONTROL INFECTIONS FOR BIAS DETECTION IN A PHASE III COVID-19 VACCINE TRIAL

Summary: *Negative control outcomes (NCOs) are useful tools for hidden bias detection, but empirical evidence validating NCOs for COVID-19 is lacking. To address this gap, we examined the blinded phase of the randomized, placebo-controlled Coronavirus Vaccine Efficacy (COVE; NCT04470427) trial of the mRNA-1273 COVID-19 vaccine. We confirmed that acute respiratory illness (ARI) with a positive test for a non-SARS-CoV-2 respiratory pathogen on a multiplex PCR panel was a valid NCO for COVID-19, considering that it was unaffected by vaccination (vaccine efficacy, VE = 3.3% (95% CI, -22.3 to 23.6)) yet strongly associated with COVID-19 (odds ratio = 2.95 (95% CI, 2.00, 4.24)). Subsequently, we leveraged non-SARS-CoV-2 pathogen-positive ARIs to detect bias in time-varying VE estimates from COVE's blinded and booster phases. Balanced incidence of non-SARS-CoV-2 pathogen-positive ARI between vaccinated and unvaccinated COVID-19-free risk sets suggested low risk of depletion of susceptibles bias in VE estimates of two-dose mRNA-1273 against COVID-19 during the blinded phase (VE = 92.5% (95% CI, 88.8, 94.9) 14 days post-dose-two, stable for 5 months). In COVE's booster phase, higher non-SARS-CoV-2 ARI incidence was observed after the single booster (intensity ratio, IR = 2.38 (95% CI, 1.75, 3.25) 14 days post-boost), suggesting that naïve booster VE estimates may underestimate VE against COVID-19. Our findings demonstrate the utility of negative control infections for unraveling complex biases in COVID-19 vaccine studies.*

2.1 Introduction

Continued SARS-CoV-2 evolution and waning COVID-19 vaccine efficacy (VE) highlight the importance of developing new vaccines and adapting existing vaccines to be durable and broadly protective against many viral strains. Randomized, double-blinded, placebo-controlled clinical trials enable reliable assessments of vaccine effects but are time-and-resource-intensive to conduct. As a consequence, randomized trials are ill-suited to rapidly updating knowledge of VE in response to

evolving pandemic conditions, such as the advent of new viral variants. Moreover, randomized vaccine trials may be infeasible to conduct due to lack of clinical equipoise and availability of vaccines off-study. Crucially, even in such "gold-standard" trials, time-varying VE estimates obtained using naïve Cox regression models are vulnerable to bias caused by differential depletion of susceptible participants (see Chapter 1.2 for more details).[1–3]

Alternatively, observational studies can provide rich information on how VE varies over time and viral variants. However, VE estimates from observational studies are especially vulnerable to bias.[4–10] Systematic differences in exposure,[4,8,11] susceptibility,[4,5] and testing behavior[12–14] between vaccinated and unvaccinated individuals can confound the relationship between vaccination and COVID-19. Misclassification of SARS-CoV-2 cases can produce bias, particularly in observational studies using test negative designs.[15–17] Traditional strategies for handling bias such as observed confounder adjustment and sensitivity analysis are limited by complex, overlapping sources of bias and the implausibility of measuring and correctly adjusting for all confounders. Therefore, alternative strategies for diagnosing bias in observational vaccines studies are needed.

Negative control outcomes (NCOs) are outcomes unaffected by the intervention but share common sources of bias with the primary outcome.[18,19] NCOs are commonly used in laboratory sciences, epidemiology, and causal inference to rule out non-causal explanations of study results. Vaccine studies have exploited negative controls in diverse ways, including to detect and reduce bias in VE estimates from observational studies[20–22] and mitigate bias due to differential outcome ascertainment in test negative designs[23–25] and cluster randomized trials.[26] To our knowledge, no studies have rigorously validated and compared NCOs for COVID-19 in a randomized, placebo-controlled trial, an important objective given their widespread use and lack of clarity around ideal choice of NCO endpoint for COVID-19.[16,22,27]

The Coronavirus Vaccine Efficacy (COVE; NCT04470427) study[28,29] was a randomized, placebo-controlled, double-blinded study of the mRNA-1273 vaccine and was an ideal case study of NCOs for COVID-19. COVE conducted symptom-triggered testing for SARS-CoV-2 and a panel of 20 other respiratory pathogens using a multiplex PCR assay. We hypothesized that acute respiratory illness (ARI) caused by "off-target" non-SARS-CoV-2 respiratory pathogen was a valid NCO for COVID-19, considering the vaccine's specificity and the potential association between non-SARS-CoV-2-pathogen positive ARIs and COVID-19 via their shared route of transmission and

overlapping risk factors. In contrast to previous studies that justified NCOs based on subject matter knowledge, COVE's unique design – a randomized, placebo-controlled trial followed by an open-label extension mimicking an observational COVID-19 vaccine study – permitted an opportunity to empirically test the NCO assumption in the blinded phase. Then, "validated" NCOs could be leveraged for downstream bias detection tasks, including detecting bias in time-varying VE estimates from COVE's blinded and booster phases.

2.2 Methods

2.2.1 COVE Study Overview

COVE was a phase 3, randomized, multi-site, double-blind study in the US adults aged $\geq$ 18 years (see Figure 2.1 for a consort diagram with study phases and pivotal dates). 30,420 participants were enrolled between July 27–October 23, 2020, and randomly assigned in a 1:1 ratio to receive two 100-μg doses of mRNA-1273 vaccine or placebo. Randomization was stratified by age <65 years, age <65 years and "high risk" of severe COVID-19, and age $\geq$65 years; this is referred to as the "randomization stratum". Participants younger than 65 were categorized as "at risk" if they had one of the following risk factors based on CDC criteria: chronic lung disease (e.g., emphysema, chronic bronchitis, idiopathic pulmonary fibrosis, cystic fibrosis, or moderate-to-severe asthma), cardiac disease (e.g., heart failure, congenital coronary artery disease, cardiomyopathy, or pulmonary hypertension), severe obesity (BMI $\geq$40), diabetes (type 1, type 2, or gestational), liver disease, or infection with human immunodeficiency virus. The period of blinded follow-up is referred to as the "blinded phase." The median duration of blinded follow-up was 4.7 months with median end date January 17, 2021. Predominantly ancestral SARS-CoV-2 circulated in the blinded phase.

Following US Food and Drug Administration Emergency Use Authorization (EUA) of mRNA-1273 in December 2020, the protocol was amended to include an open-label phase where participants were offered to be unblinded and placebo recipients were offered two doses of mRNA-1273. Among participants who began the open-label phase, median follow-up was 14.3 months after unblinding. In the open-label phase, early SARS-CoV-2 variants emerged including Alpha followed by Delta. On September 23, 2021, the protocol was amended again to offer a single 50-μg mRNA-1273 booster to

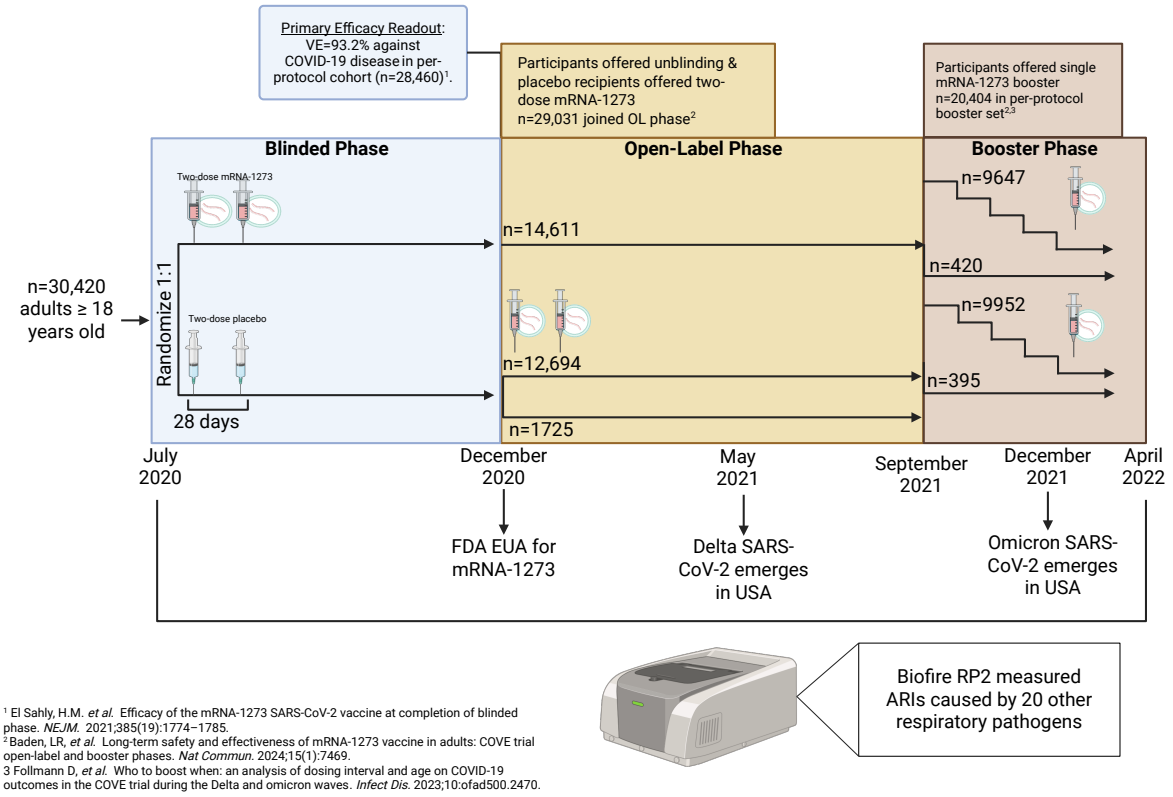

Figure 2.1: Illustrative consort diagram displaying COVE study phases alongside major COVID-19 pandemic hallmark dates.

vaccinated participants. The period from September 23, 2021, until the data-cutoff date of April 5, 2022, is referred to as the "booster phase." The booster phase emulated a prospective observational COVID-19 vaccine study, as participants elected whether and when to receive the booster. Median followup post-boost was 5.3 months. In the booster phase, the SARS-CoV-2 lineages were Delta followed by Omicron.

Respiratory Infection Surveillance

Throughout COVE, participants with respiratory illness symptoms lasting ≥48 hours were referred to the clinical site within 72 hours or visited at home for an "illness visit." At the illness visit, two nasopharyngeal (NP) swabs were collected. One was tested for SARS-CoV-2 using RT-PCR and the other was tested for 20 other respiratory pathogens using a multiplex PCR assay (Biofire RP2,

Eurofins Viracor). If a clinic or home visit was not possible, participants submitted a saliva sample or nasal swab for SARS-CoV2 testing. Whole genomic sequencing was performed on swabs positive for SARS-CoV-2 to identify the viral strain.[27]

Endpoints

A COVID-19 endpoint is defined as first occurrence of COVID-19 at least 14 days post-dose-two of mRNA-1273/placebo. This was the COVE trial's primary endpoint.[28,29] In the booster phase, a COVID-19 endpoint was defined as the first COVID-19 occurrence before or at least 14 days after booster. All COVID-19 endpoints required at $\geq$1 NP/nasal swab or saliva sample positive for SARS-CoV-2 by RT-PCR, and $\geq$ 2 systemic symptoms or $\geq$ 1 respiratory symptom (see Section 1.3, Appendix S1).[28,29] The COVID-19 event time is date of first positive SARS-CoV-2 result by RT-PCR after symptom onset.

We examined two candidate NCO endpoints. The first, non-SARS-CoV-2 pathogen-positive ARI, was a respiratory illness visit with a positive test for any non-SARS-CoV-2 pathogen. No clinical case definition was used; however, participants were only tested when experiencing respiratory illness symptoms. Coinfections with SARS-CoV-2 were counted. The event time is the date of positive test result for a non-SARS-CoV-2 pathogen. Informed by the widespread use of test negative designs,[23,24,30] the second candidate NCO was SARS-CoV-2-negative ARI, defined as an illness visit with symptoms matching the COVID-19 case definition where the first swab or saliva sample collected within 3 days of symptom onset was negative for SARS-CoV-2 by RT-PCR. This endpoint includes non-SARS-CoV-2 pathogen-positive ARIs and illness visits negative for all pathogens. The event time is date of first SARS-CoV-2-negative ARI.

Analysis Cohorts

The full analysis set (FAS) included all randomized participants who received at least one dose of mRNA-1273 or placebo (n=30,347). The FAS is the basis for descriptive analysis. The blinded phase per-protocol cohort consisted of FAS participants who received two doses of mRNA-1273 or placebo with no major protocol deviations and no immunologic/virologic evidence of prior SARS-CoV-2 infection at enrollment (n = 28,460). The booster phase per-protocol cohort consisted

of participants who previously received the two-dose mRNA-1273 primary series without major protocol deviations and without virological/immunological evidence of prior SARS-CoV-2 infection on September 23, 2021 (n=20,404).[31]

2.2.2 Bias detection methodology

Below, we discuss an approach to detect confounding and selection bias in hazard-based estimates of time-varying vaccine effectiveness using negative control infection outcomes. For simplicity, we appeal to a randomized vaccine trial, where confounding is absent by design but selection bias caused by differential rates of depletion of susceptible participants from the vaccine and placebo arms can distort estimates of waning VE (see Chapter 1.2 for an overview).

Consider the time-varying hazard ratio, which by definition, contrasts the instantaneous infection rates t-days post-baseline between vaccinated and unvaccinated individuals:

$$\text{HR}(t) = \lim_{\epsilon \downarrow 0} \frac{P(t < Y \leq t + \epsilon | A = 1, Y \geq t)}{P(t < Y \leq t + \epsilon | A = 0, Y \geq t)}.$$

Let $Y(a)$ denote the potential primary infection time under hypothetical assignment to treatment $A = a$. By virtue of randomization, exchangeability holds at baseline: $\{Y(1), Y(0)\} \perp A$. Under exchangeability and consistency of potential outcomes, we rewrite the time-varying hazard ratio as follows:[32]

$$\text{HR}(t) = \lim_{\epsilon \downarrow 0} \frac{P(t < Y(1) \leq t + \epsilon | Y(1) \geq t)}{P(t < Y(0) \leq t + \epsilon | Y(0) \geq t)}.$$

Therefore, the time-varying hazard ratio mixes the vaccine effect at time t with a comparison of two different groups (i.e., principal strata) of individuals: vaccinated survivors $\{Y(1) \geq t\}$ and unvaccinated survivors $\{Y(0) \geq t\}$. Hence, $\text{HR}(t) \neq 1$ may be due to the vaccine effect at time t, differences between the vaccinated and unvaccinated survivors at time t, or a combination of both.

To illustrate why the time-varying hazard ratio is commonly considered "biased" for any meaningful notion of time-varying VE, we introduce a frailty variable U which summarizes a possibly large number of unmeasured variables influencing infection risk collected just prior to time t. Suppose we assume latent independence of $Y(1)$ and $Y(0)$ conditional on the frailty U.

Assumption 1 (Latent potential outcome independence)**.** *Assuming U is sufficient to capture all sources of heterogeneity in primary infection risk, such that $Y(1) \perp Y(0) \mid U$.*

Under Assumption 1, the time-varying hazard ratio can be written as follows using Tower Law.

$$\mathrm{HR}(t) \overset{\mathrm{A1}}{=} \lim_{\epsilon \downarrow 0} \frac{\int P(t < Y(1) \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u | Y(1) \geq t) du}{\int P(t < Y(0) \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u | Y(0) \geq t) du}$$

Despite the fact that randomization ensures exchangeable distributions of U between vaccine and placebo groups at baseline, the time-varying hazard ratio is biased for a principal stratum causal effect[33] due to marginalization over different distributions of U. When a vaccine is protective, "high-risk" participants will succumb to infection more rapidly in the placebo group, leading to frailty distributions that become increasingly different over time: $f_U(u \mid Y(1) \geq t) \overset{d}{\neq} f_U(u \mid Y(0) \geq t)$.

Our task is to devise a method to detect bias in $\mathrm{HR}(t)$. Suppose the randomized trial measures an auxiliary outcome N that is a negative control outcome, satisfying the following assumption.

Assumption 2 (NCO)**.** *N is a valid NCO: $N = N(1) = N(0)$.*

We introduce an alternative parameter:

$$\Psi_N(t) := \lim_{\epsilon \downarrow 0} \frac{P(t < N \leq t + \epsilon | Y \geq t, A = 1)}{P(t < N \leq t + \epsilon | Y \geq t, A = 0)} = \lim_{\epsilon \downarrow 0} \frac{P(t < N \leq t + \epsilon | Y(1) \geq t)}{P(t < N \leq t + \epsilon | Y(0) \geq t)}.$$

$\Psi_N(t)$ uses the same conditioning events ($\{Y(1) \geq t\}$ and $\{Y(0) \geq t\}$) as the numerator and denominator of the time-varying hazard ratio, preserving the structural source of selection bias in the time-varying hazard ratio.[32] However, $\Psi_N(t)$ swaps out the primary outcome for the negative control outcome (Assumption 2) in the probability statement, which eliminates the potential for a causal effect at time t. Using the NCO in this way is similar to the "leave-out-an-essential-ingredient" technique used in experimental science;[18] replacing the primary outcome with the negative control outcome eliminates a causal vaccine effect and isolates the impact of systematic differences (i.e., bias) between vaccinated and unvaccinated survivors. Hence, under Assumption 2, $\Psi_N(t) \neq 1$ can be used to diagnose bias in time-varying hazard ratios.

While $\Psi_N(t) \neq 1$ is sufficient evidence of depletion-of-susceptibles bias, the converse need not hold; bias may still be present even when $\Psi_N(t) = 1$. For $\Psi_N(t)$ to serve as a reliable diagnostic

for depletion-of-susceptibles bias, the primary outcome Y and negative control outcome N must share the same underlying sources of unmeasured heterogeneity. We formalize this requirement in Assumption 3 below.

Assumption 3 (U-comparability[18,19]). *Assume N and Y depend on the same unmeasured frailty variable such that $N \not\perp U$ and $N \perp (Y(1), Y(0))|U$.*

Under Assumption 3, $\Psi_N(t)$ can be used to quantify the imbalance in frailty distributions that generates bias in the time-varying hazard ratio.

$$\Psi_N(t) \overset{(A3)}{=} \lim_{\epsilon \downarrow 0} \frac{\int P(t < N \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u|Y(1) \geq t) du}{\int P(t < N \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u|Y(0) \geq t) du}$$

We briefly consider the case where vaccination (A) is not randomized, such as in an observational study. Now, time-varying VE estimates may be distorted due to confounding in addition to depletion of susceptibles bias. However, the logic of our approach is mostly unchanged, as confounding and depletion of susceptibles are merely mechanisms that lead to lack of exchangeability between vaccinated and unvaccinated disease-free cohorts. We assume latent exchangeability, which essentially presumes that U captures all unmeasured confounders.

Assumption 4 (Latent exchangeability). *Assume $\{N, Y(1), Y(0)\} \perp A \mid U$ holds. This is equivalent to saying that participants are essentially randomized within levels of the frailty.*

Under Assumption 1 to 4, Ψ_N can be written as,

$$\Psi_N(t) = \lim_{\epsilon \downarrow 0} \frac{\int P(t < N \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u|Y(1) \geq t) du}{\int P(t < N \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u|Y(0) \geq t) du}$$

which will capture the bias in the time-varying hazard ratio due to selection and/or confounding by U.

An important secondary question is under what conditions $\Psi_N(t)$ may be used to *quantify* the bias caused by the frailty distribution imbalance. To proceed, we decompose the time-varying hazard ratio into the product of a meaningful causal effect and a bias term.

$$\mathrm{HR}(t) \overset{\mathrm{A1}}{=} \lim_{\epsilon \downarrow 0} \frac{\int P(t < Y(1) \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u|Y(1) \geq t) du}{\int P(t < Y(0) \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u|Y(1) \geq t) du}$$

$$\times \left(\frac{\int P(t < Y(0) \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u | Y(1) \geq t) du}{\int P(t < Y(0) \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u | Y(0) \geq t) du} \right)$$

The first term on the right-hand-side can be interpreted as the time-varying principle stratum vaccine effectiveness standardized to the frailty distribution of vaccinated survivors. The latter term is a bias term that quantifies the drift in vaccine-free potential outcomes resulting from marginalizing over different frailty distributions.

In order to accurately quantify the bias using $\Psi_N(t)$, a sufficient condition is that the relative rates of failure from the vaccine-preventable and negative control infections are stable/proportional across levels of the frailty.

$$\gamma(t) = \lim_{\epsilon \downarrow 0} \frac{P(t < Y(0) \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u)}{P(t < N \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u)} \qquad \forall t > 0, u \text{ with positive support}$$

Under this condition, $\Psi_N(t)$ quantifies the bias precisely.

$$\begin{aligned} \Psi_N(t) &= \lim_{\epsilon \downarrow 0} \frac{\int P(t < N \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u | Y(1) \geq t) du}{\int P(t < N \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u | Y(0) \geq t) du} \\ &= \lim_{\epsilon \downarrow 0} \frac{\cancel{1/\gamma(t)} \int P(t < Y(0) \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u | Y(1) \geq t) du}{\cancel{1/\gamma(t)} \int P(t < Y(0) \leq t + \epsilon | Y(0) \geq t, Y(1) \geq t, U = u) f_U(u | Y(0) \geq t) du} \end{aligned}$$

2.2.3 Statistical Analysis

Analyses described herein are post-hoc and may differ from protocol-specified primary analyses.

We characterized SARS-CoV-2 and non-SARS-CoV-2 pathogen-positive ARIs over study time, calendar time, randomization arm, randomization strata, and baseline characteristics using case counts and incidence rates estimated using exact Poisson methods.

We tested the validity of the two candidate NCOs by calculating treatment-arm-specific covariate-adjusted cumulative incidence curves among the blinded phase per-protocol cohort.[34] VE against each endpoint was computed as one minus the vaccine:placebo ratio of arm-specific cumulative incidences at 180 days post-randomization. All estimators adjusted for sex, racial/ethnic under-represented minority (URM) status, randomization stratum, and continuous baseline risk score[35] to improve precision and ensure validity under random censoring conditional on covariates. A Cochran–Mantel–Haenszel odds ratio stratified by randomization arm summarized the association

between non-SARS-CoV-2 pathogen-positive ARI and COVID-19 statuses in the blinded phase.

We estimated time-varying VE against COVID-19 for the two-dose primary series and single booster using proportional hazards models on calendar time scales. We used a calendar time scale to avoid violations of proportional hazards caused by variation in disease incidence over calendar time.[36] $t = 0$ referred to the date of enrollment of the first study participant. The hazard model for vaccine-preventable pathogen (COVID-19) and intensity model for the negative control (non-SARS-CoV-2 pathogen-positive ARIs) on date t are shown in 2.1 below.

$$\begin{aligned} h_1(t \mid V, E) &= I(E \geq t)(1 - I(0 \leq t - V < B))h_{01}^S(t) \exp\left(I(V \leq t)\left\{\beta_0^1 + \sum_{\ell=1}^{L} \beta_\ell^1 P_\ell(t - V; \mathbf{k}, \delta)\right\} + X\theta^1\right) \\ h_0(t \mid V, E) &= I(E \geq t)(1 - I(0 \leq t - V < B))h_{00}^S(t) \exp\left(I(V \leq t)\left\{\beta_0^0 + \sum_{\ell=1}^{L} \beta_\ell^0 P_\ell(t - V; \mathbf{k}, \delta)\right\} + X\theta^0\right) \end{aligned} \tag{2.1}$$

Where V denotes the date of first dose/booster vaccination date, E denotes the date of study entry, S refers to a unique combination of sex, racial/ethnic underrepresented minority (URM) status, and randomization stratum, X denotes the continuous baseline risk score with coefficients θ^1 and θ^0,[35] and $\{P_\ell\}_{\ell=1}^L$ denotes a P-spline basis of dimension L, with degree δ, vector of knot locations $\mathbf{k}$, and basis coefficient vectors $\boldsymbol{\beta}^1$ and $\boldsymbol{\beta}^0$. The degrees of freedom were chosen by maximizing a corrected AIC.[37,38] Following common practice in vaccine studies and to harmonize with prior COVE study analyses, we avoided counting infections and person-time in a "blackout period" from $[0, B)$ days after the start of the vaccination series to better capture the full benefit of immunization.[28,29,31,36,39] For the two-dose primary series and booster analyses, B was chosen to be 41 and 13 respectively, meaning that the effect of the vaccine was estimated starting 14 days after the completion of the primary series/single booster respectively, the presumed "peak" of vaccine protection. VE against SARS-CoV-2 and intensity ratio (IR) of negative control infections were calculated as follows

$$\widehat{\mathrm{VE}}(s) = 1 - \exp\left(\hat{\beta}_0^1 + \sum_{\ell}^{L} \hat{\beta}_\ell^1 P_\ell(s; \mathbf{k}, \delta)\right), \qquad \widehat{\Psi}_N(s) = \exp\left(\hat{\beta}_0^0 + \sum_{\ell}^{L} \hat{\beta}_\ell^0 P_\ell(s; \mathbf{k}, \delta)\right).$$

All outcomes were censored at earliest event of discontinuation, last event assessment, study termination, off-study COVID-19 vaccination, and the data cutoff date. Participants were also censored at SARS-CoV-2 infections occurring during the blackout period. The blinded phase analysis

also censored participants at the earliest of the events described above and their unblinding date.

In the booster phase analysis, a common time origin of $E = 0$ corresponding to September 23, 2021 (the start of the booster phase) was used for all participants. To distinguish between waning vaccine effectiveness from changes in circulating SARS-CoV-2 variants over calendar time in the booster phase, we estimated VE against Delta and Omicron variants separately using the Prentice et al. (1978) competing risk framework. Missing SARS-CoV-2 lineages were singly imputed using per-week lineage frequencies in the US based on GISAID estimates.[41] Since Delta quickly and completely transitioned to Omicron SARS-CoV-2 in the Winter 2021 in United States, we believed uncertainty in variant imputations to be very small (see Figure 2.2).

To detect depletion of susceptibles in time-varying VE estimates against COVID-19 for the two-dose primary series and single-dose booster, we estimated the bias detection parameter $\Psi_N(t-v)$ using the intensity ratio based on the semiparametric model for the intensity shown in 2.1. Non-SARS-CoV-2 pathogen-positive ARIs were censored at the date of COVID-19 infection to align the risk sets and thereby ensure the structural source of bias in the primary analysis was preserved in the negative control analysis.[32]

A supporting analysis estimated the two-dose mRNA-1273 VE over the blinded and open-label phases post-crossover following the methods of Fintzi and Follmann[36]. A parallel analysis of negative control infections was used to detect potential selection bias and confounding. Results can be found in Appendix A.3.2. The method assumes that the efficacy profile of the vaccine during the blinded phase applied to open-label phase, which was likely violated due to the advent of Delta SARS-CoV-2 post-crossover.

All analyses were conducted using R version 4.2.3 (R Core Team, Vienna, Austria) based on dataset with cutoff date of April 5, 2022.

2.3 Results

2.3.1 Circulating Respiratory Infections in COVE

Incident respiratory infections throughout COVE are shown in Figure 2.2. SARS-CoV-2 cases in COVE reflected COVID-19 trends observed in the broader US population. Human rhinovirus/enterovirus was the dominant non-SARS-CoV-2 pathogen-positive ARI early in COVE. Easing pandemic re-

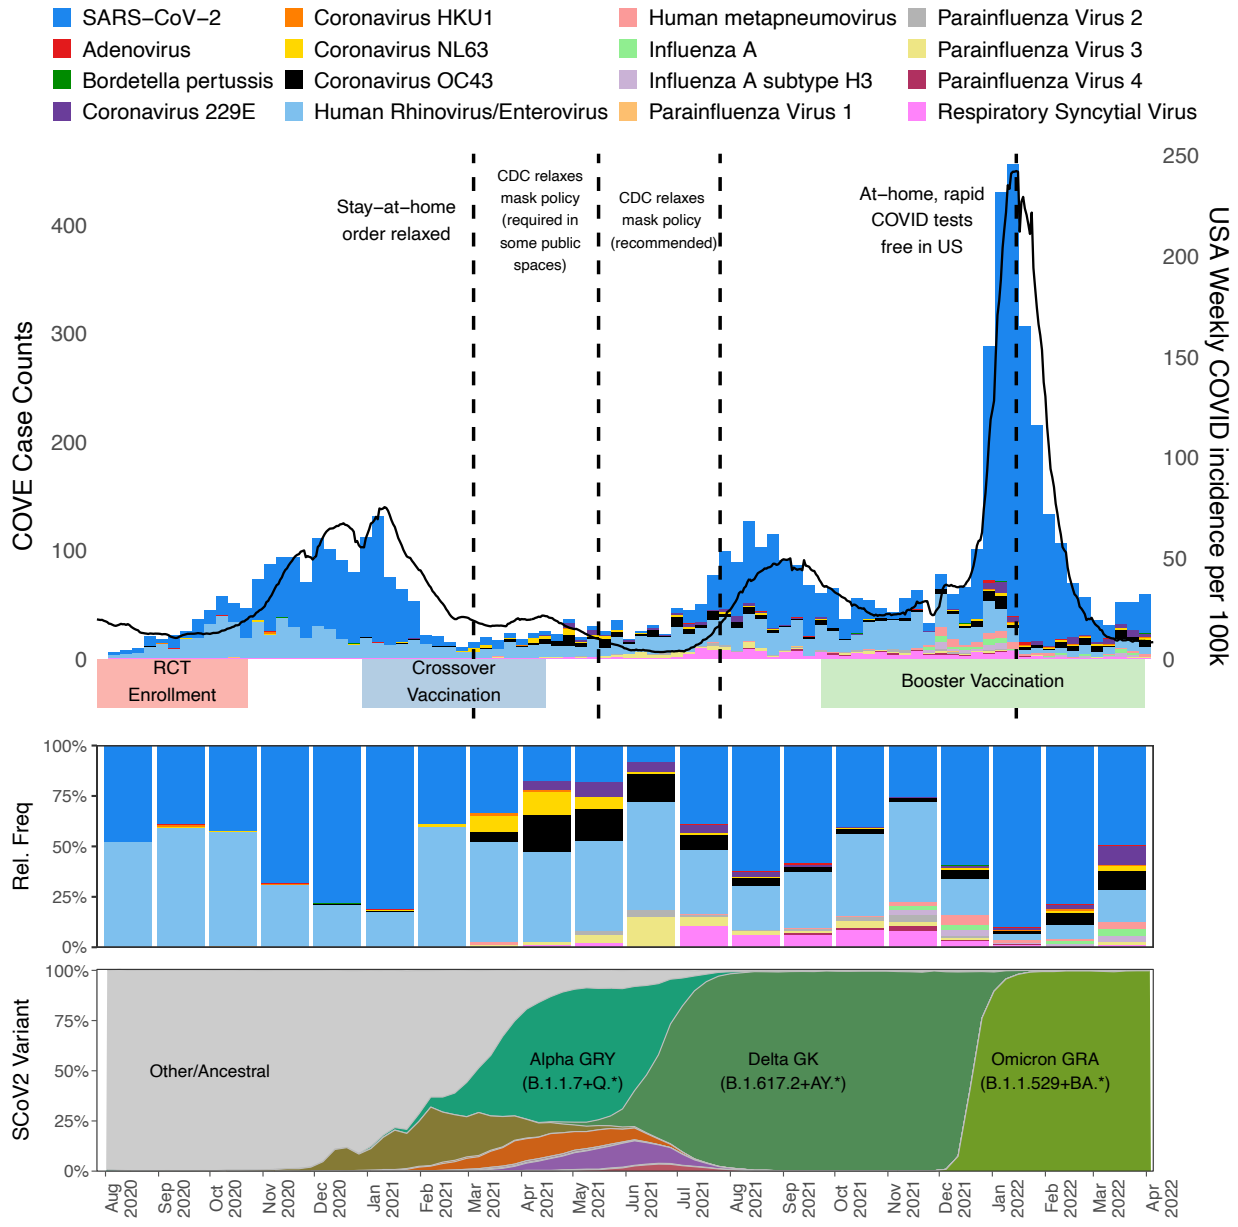

Figure 2.2: Number of incident respiratory infections over the duration of the COVE study in the FAS ($n = 30,347$). Shown in black is the smoothed COVID-19 incidence rate per 100,000 individuals in the US population derived from Center for Systems Science and Engineering (CSSE) at Johns Hopkins University.[42] Dotted vertical lines correspond to pandemic hallmark dates including relaxations of stay-at-home orders, mask recommendations, and when at-home COVID-19 tests were offered for free in the United States. Shown beneath are relative frequencies of respiratory infections per month and composition of SARS-CoV-2 variants in the United States over time based on GISAID data.[41]

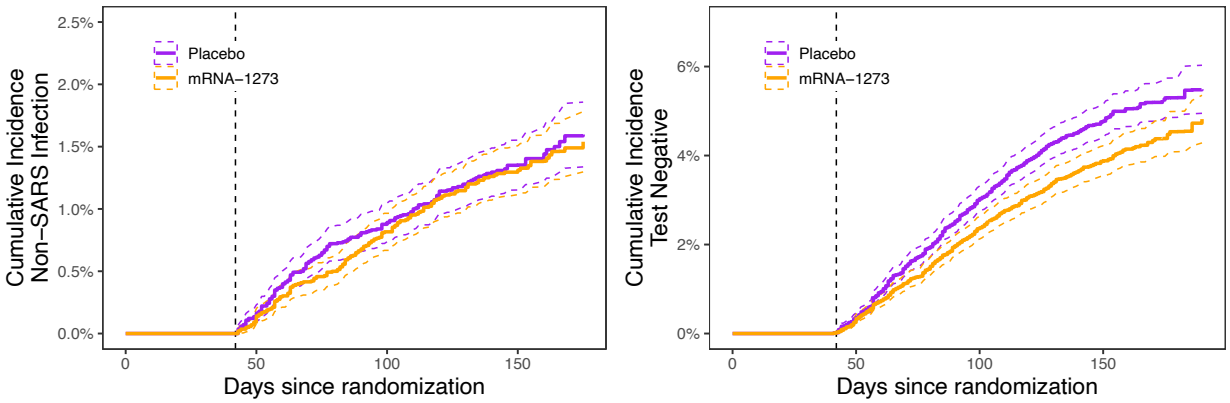

Figure 2.3: Cumulative incidence of non-SARS-CoV-2 pathogen-positive ARI (left) and SARS-CoV-2-negative ARI (right) counted 14 days post dose two in the blinded phase. Curves were estimated from the blinded phase per-protocol population. Cumulative incidence estimates adjusted for sex, racial/ethnic URM status, randomization strata, and baseline risk score. 95% pointwise confidence intervals shown.

strictions was associated with a diversification of respiratory infections including Coronaviruses NL63, OC43, and 229E, followed by Parainfluenza virus types 2 and 3 and respiratory syncytial virus (RSV) in May–July 2021 and influenza starting in November 2021. Availability of at-home COVID-19 tests and the Omicron wave in January 2022 coincided with detection of fewer non-SARS-CoV-2 pathogen-positive ARIs.

Lower incidence of non-SARS-CoV-2 pathogen-positive ARIs was associated with age $\geq$65 years and male sex (blinded phase incidence rate ratio, IRR = 0.428 (95% CI, 0.329, 0.558); IRR = 0.739 (95% CI, 0.621, 0.879)). A marginal association between underrepresented racial/ethnic minorities and higher rates of non-SARS-CoV-2 pathogen-positive ARI was also observed (see Appendix A.3.1 for results).

2.3.2 No effect of mRNA-1273 on non-SARS-CoV-2 pathogen-positive ARI in COVE blinded phase

To assess if non-SARS-CoV-2 pathogen-positive ARI and SARS-CoV-2-negative ARI endpoints were valid NCOs, we compared the cumulative incidences of each endpoint across randomization arms in COVE's blinded phase (Figure 2). mRNA-1273 vaccination had no apparent effect on non-

SARS-CoV-2 pathogen-positive ARI (VE = 3.3% [95% CI, -22.3 to 23.6]). In contrast, mRNA-1273 vaccination significantly reduced the rate of SARS-CoV-2-negative ARI (VE = 14.4% [95% CI, 2.5, 24.8]). Results were similar when outcomes were counted immediately after randomization (see Appendix A.3.2).

2.3.3 Non-SARS-CoV-2 pathogen-positive ARI is a proxy for higher COVID-19 exposure in COVE blinded phase

To evaluate if non-SARS-CoV-2 pathogen-positive ARI and COVID-19 were associated through unmeasured common causes, we examined the association between non-SARS-CoV-2 pathogen-positive ARI and COVID-19 in COVE's blinded phase. Non-SARS-CoV-2 pathogen-positive ARI was associated with an estimated 3-fold higher odds of experiencing COVID-19 (conditional OR = 2.95 [95% CI, 2.00, 4.24]). Placebo recipients who experienced a non-SARS-CoV-2 pathogen-positive ARI had a 180-day cumulative COVID-19 incidence of 18.0% (95% CI, 11.4, 24.5) compared to 7.0% (95% CI, 6.4, 7.5) for placebo recipients who avoided non-SARS-CoV-2 pathogen-positive ARI. We conclude that non-SARS-CoV-2 pathogen-positive ARI as a strongly prognostic marker of higher likelihood of COVID-19, presumably through shared unmeasured risk factors (see Appendix A.2 and A.3.2).

2.3.4 No evidence of depletion of susceptibles bias in time-varying VE estimates of two-dose mRNA-1273 during blinded phase

Time-varying VE estimates of two-dose mRNA-1273 against COVID-19 in the blinded phase are shown in Figure 2.4. Estimated VE 14 days post-second dose was 92.5% (95% CI, 88.7, 95.0) and remained stable for up to 150 days. The IR of non-SARS-CoV-2 pathogen-positive ARI was almost exactly one, suggesting that vaccinated and placebo risk sets were exchangeable in their average susceptibility of off-target infections. Hence, our analysis of the non-SARS-CoV-2 pathogen-positive ARIs did not yield evidence of systematic differences between vaccinated and unvaccinated COVID-19-free participants caused by differential depletion of susceptible participants.

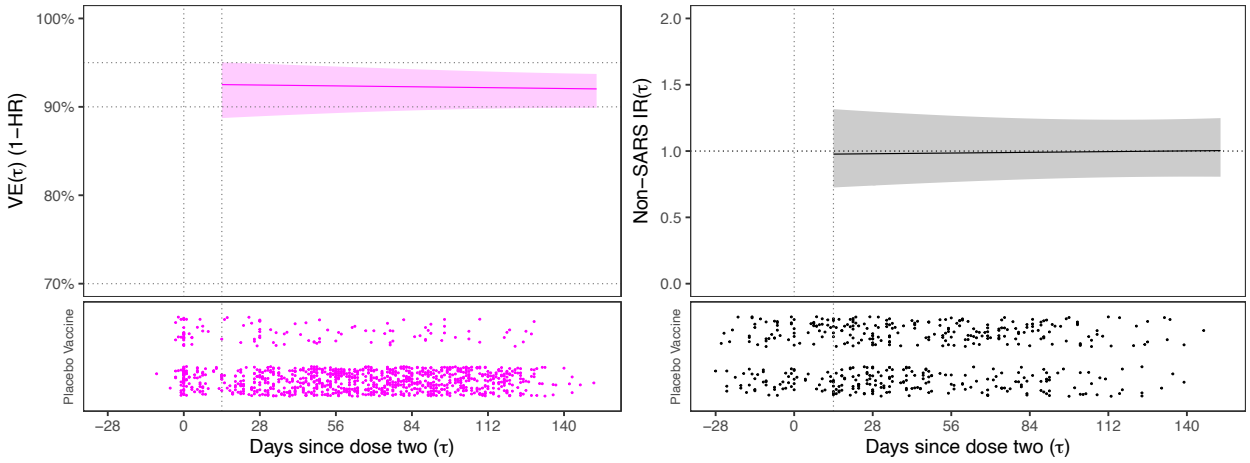

Figure 2.4: Time-varying VE of two-dose mRNA-1273 against COVID-19 (left) and IR of non-SARS-CoV-2 pathogen-positive ARI (right) as a function of days since dose two. 95% Wald CIs shown. Dots in bottom panels denote COVID-19 and non-SARS-CoV-2 pathogen-positive ARI events grouped by randomization arm. See 2.2.3 for additional details on estimation.

2.3.5 Evidence of downward bias in time-varying VE estimates of single mRNA-1273 booster

Time-varying VE estimates of the single mRNA-1273 booster are presented in Figure 2.5. Estimated VE against Delta COVID-19 was 87.5% (95% CI, 56.2.5-96.4) 14 days post-boost and appeared stable for 85 days until Delta circulation stopped. Estimated VE against Omicron COVID-19 was 58.1% (95% CI, 47.3, 66.7) 14 days post-boost and decayed over 120 days (VE = 12.2% [95% CI, -4.9 to 26.5]). The rate of non-SARS-CoV-2 pathogen-positive ARI was 2.38-fold (95% CI, 1.75, 3.25) higher 14-days after boosting compared to unboosted participants on the same calendar date. The IR declined at 60 days post-boost (IR = 1.35 [95% CI, 1.01, 1.80]) but was higher at 120 days (IR = 1.60, [95% CI, 1.05, 2.45]). The higher rate of non-SARS-CoV-2 pathogen-positive ARI observed post-boost suggested that persistent, systematic differences between boosted and unboosted participants could have led to underestimation of time-varying VE estimates against COVID-19.

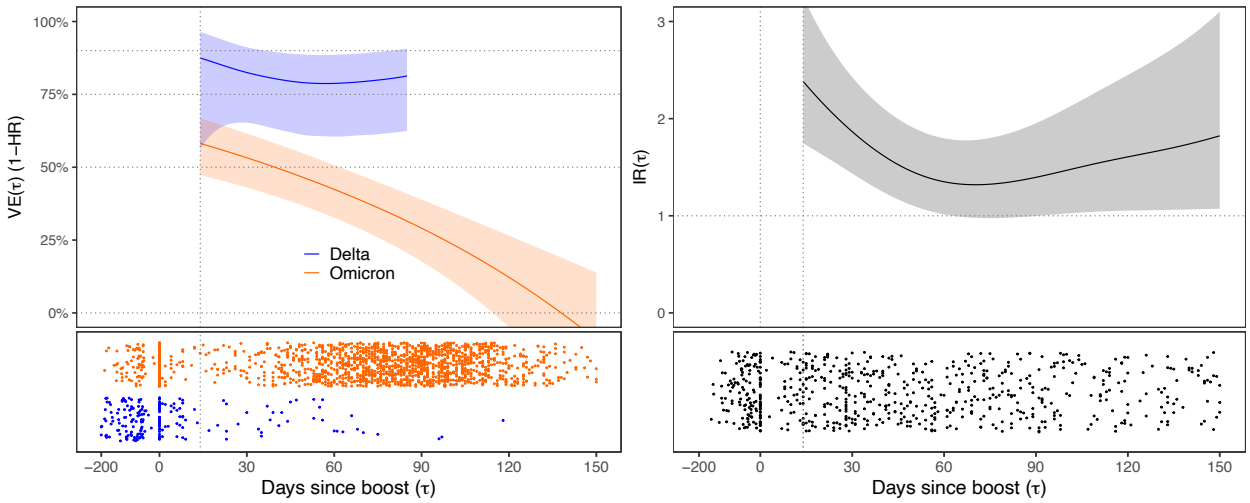

Figure 2.5: Time-varying VE of two-dose mRNA-1273 against Delta and Omicron COVID-19 (left) and IR of non-SARS-CoV-2 pathogen-positive ARI (right) as a function of days since booster. 95% Wald CIs shown. Dots with negative days since vaccination correspond to infections before boosting. Dots aligned at zero correspond to infections in individuals who were never boosted. See 2.2.3 for additional details on estimation.

2.4 Discussion

We described circulating respiratory infections in the COVE study, through periods with varying non-pharmaceutical interventions (NPI). In the blinded phase, non-SARS-CoV-2 pathogen-positive ARI was confirmed as a promising negative control outcome (NCO) for COVID-19, considering it was unaffected by mRNA-1273 vaccination yet strongly associated with COVID-19. We leveraged non-SARS-CoV-2 pathogen-positive ARIs to detect bias in time-varying VE estimates in COVE's blinded and booster phases. While we observed no evidence of depletion of susceptibles bias in two-dose mRNA-1273 VE estimates in COVE's blinded phase, we observed higher incidence of non-SARS-CoV-2 pathogen-positive ARI after the single mRNA-1273 booster, which suggests that booster VE estimates may be biased underestimates.

Consistent with previous reports,[43] human rhinovirus/enterovirus was the main circulating non-SARS-CoV-2 respiratory pathogen in COVE from July 2020–March 2021. Persistence of rhinovirus despite strict NPI was documented across geographically diverse settings,[44–51] potentially due to

its lack of viral envelope, stability on surfaces, and genomic diversity.[52] Other non-SARS-CoV-2 respiratory pathogens were largely absent, agreeing with previous reports of reduced circulation of influenza,[44,47,53–56] RSV,[47,55,57–60] seasonal coronaviruses,[44,47,61,62] parainfluenza virus,[47,48,61,63], and metapneumovirus.[47,48,61,63] As NPI relaxed, respiratory pathogens diversified, consistent with other studies.[64,65] Fewer non-SARS-CoV-2 pathogen-positive ARIs were detected post-January 2022, which may reflect lower ascertainment due to available at-home COVID-19 testing. We observed lower incidence of non-SARS-CoV-2 pathogen-positive ARI among older male participants, perhaps reflecting differences in NPI adherence or occupational exposure to respiratory pathogens.[66]

To detect bias, a valid NCO must be unaffected by the intervention and share unmeasured causes with the outcome of interest.[19] Usually, a candidate NCO is justified using subject matter knowledge. However, the blinded phase of COVE offered a unique opportunity to empirically test the NCO assumption for two candidate NCO endpoints. We identified that acute respiratory illnesses accompanied by a positive PCR test for a non-SARS-CoV-2 respiratory pathogen were excellent NCOs; they were not affected by mRNA-1273 vaccination yet positively correlated with COVID-19 outcomes, presumably through unmeasured common causes. We cannot confirm if non-SARS-CoV-2 pathogen-positive ARI is a valid NCO for other COVID-19 vaccine constructs, although, we hypothesize that the NCO assumption will hold for similarly antigen-specific vaccines (i.e., mRNA, DNA, protein subunit, or viral vector vaccines).

In contrast, we observed that the other candidate NCO, SARS-CoV-2-negative ARIs, were significantly affected by COVID-19 vaccination in COVE's blinded phase. One explanation for this result is the vaccine cross-protected against a fairly common non-SARS-CoV-2 respiratory pathogen that was not included in the multiplex panel. While cross-protection cannot be discounted, we consider it unlikely due to the vaccine's specificity and the absence of other coronaviruses during COVE's blinded phase. Another possibility is incomplete blinding (perhaps due to vaccine reactogenicity) led to differences in participant behavior between treatment arms. However, balanced non-SARS-CoV-2 pathogen-positive ARI rates supports that blinding was maintained. In our view, the most plausible explanation is outcome misclassification, meaning the "SARS-CoV-2 test-negative ARI" case definition incorrectly classified some true SARS-CoV-2 infections as test-negative illnesses, despite those infections being preventable by vaccination. Misclassification is supported by evidence that mRNA-1273 vaccination reduces SARS-CoV-2 viral loads, which could reduce detection of

SARS-CoV-2 infections.[43,67] While the exact mechanism remains unclear, our findings suggest that ARIs accompanied by a positive test for an off-target pathogen may be a superior choice of negative control than ARIs with a negative test for the target pathogen. This insight is relevant for observational studies using test negative designs (TNDs). To reliably estimate vaccine effects in a TND, typically one must assume vaccination does not affect control infections.[19,68–70] Our results support that non-SARS-CoV-2 pathogen-positive ARI better satisfies the negative control assumption and may capture sources of confounding beyond test-seeking behavior. For example, correlated uptake of vaccines against COVID-19 and other vaccine-preventable respiratory diseases can bias VE estimates in TNDs using a test negative controls.[68] Multiplex pathogen testing enables the selection of alternative control endpoint-—like ARI caused by human rhinovirus—–that can avoid this bias. In routine care settings where TNDs are frequently conducted, multiplex pathogen testing may be infeasible, and test negative ARI may be an acceptable alternative.

By leveraging non-SARS-CoV-2 pathogen-positive ARI as a negative control outcome (NCO), we found no evidence that differential depletion of susceptibles biased time-varying VE estimates of two-dose mRNA-1273 during the blinded phase of COVE. Although negative controls are commonly used to detect unmeasured confounding, their utility for detecting selection bias has received comparatively less attention.[71] Our bias detection approach is conceptually related to methods proposed by Lipsitch, Ray, and colleagues for evaluating waning influenza vaccine effectiveness.[2,72] Their strategy reduced depletion-of-susceptibles bias by vaccinating all participants before the flu season began, thereby temporarily eliminating exposure, an "essential ingredient" for bias to manifest. Our approach detects depletion-of-susceptibles bias by applying the same selection criteria built into hazard-based VE analyses to off-target infection endpoints. Under the negative control assumption, any time-varying association between vaccination and off-target infection reflects selection-driven, systematic differences between vaccinated and unvaccinated COVID-19-free risk sets that the hazard-based, primary analysis describes. The sensitivity of our approach depends on the incidence of off-target infections and the extent to which COVID-19 and off-target infections share common unmeasured risk factors. One important limitation is that undetected immunizing SARS-CoV-2 infections are not captured by the negative control outcome because SARS-CoV-2 immunity does not directly influence susceptibility to off-target infections. To mitigate this issue, the COVE analyses excluded individuals with prior virological or serological evidence of SARS-

CoV-2 infection. Finally, we emphasize that the two-dose VE estimates considered here reflect a relatively short follow-up period (median 4.7 months) during circulation of predominantly ancestral SARS-CoV-2 strains.

In the booster phase, we identified higher non-SARS-CoV-2 pathogen-positive ARI rates after booster vaccination. One hypothesis explaining this result is negative viral interference,[73] where mRNA-1273 indirectly affected non-SARS-CoV-2 pathogen-positive ARIs by delaying the occurrence of COVID-19. Although this hypothesis cannot be discounted, the positive association between COVID-19 and non-SARS-CoV-2 pathogen-positive ARIs in COVE's blinded phase and the absence of observed interference during prior COVID-19 waves disputes this hypothesis. We believe a more likely explanation is that imbalance in non-SARS-CoV-2 pathogen-positive ARI rates reflects lack of exchangeability between boosted and unboosted cohorts. As suggested by the blinded phase analysis, if higher non-SARS-CoV-2 pathogen-positive ARIs indicates higher SARS-CoV-2 exposure, our booster VE estimates would underestimate the true VE that would have been observed in a target trial where booster status was randomized. The time-varying association between booster and non-SARS-CoV-2 pathogen-positive ARI suggest that early (<30 days post-boost) and late (>100 days post-boost) efficacy estimates were subject to the greatest degree of bias.

One limitation of NCOs is they cannot *distinguish* between different sources of bias.[18] There were myriad potential sources of confounding in COVE's booster phase, including baseline differences in boosted and unboosted participants,[4,5,8] differences between participants boosted on different calendar dates, and time-varying differences such as changes in exposure resulting from behavioral disinhibition.[4,5] Selection bias caused by differential depletion of susceptible participants was also possible and is supported by higher non-SARS-CoV-2 incidence observed >100 days post-boost, as selection effects should be largest over long periods of follow-up with high infection burdens. These biases are not mutually exclusive and may have combined to distort booster VE estimates. Another limitation is that many pathogens circulating in the booster phase were absent in the blinded phase, likely due to strict NPI. Hence, we could not test the NCO assumption for individual non-SARS-CoV-2 pathogens absent in the blinded phase, such as influenza and RSV. However, the blinded phase analysis supports the specificity of mRNA-1273's effect, and the higher non-SARS-CoV-2 pathogen-positive ARI incidence observed post-boost does not support the hypothesis of cross-protection. Future studies conducted under less strict NPI could test the NCO hypothesis against

individual non-SARS-CoV-2 respiratory pathogens.

Validating negative control endpoints in late-phase trials can inform the design and analysis of future observational vaccine studies. The US hospital reporting requirements for SARS-CoV-2, influenza, and RSV infections could facilitate the application of negative control approaches in real-world effectiveness studies of COVID-19 vaccines using electronic health records.[74] Developing statistical approaches to remove bias from time-varying VE estimates using NCOs[75–77] or double-negative controls[78–80] is a compelling avenue for future work. Moreover, NCOs could be useful for bias detection and elimination in other non-randomized analyses of vaccines such as immune correlates analyses.[35,81] We hope this work will motivate creative uses of negative controls in vaccine studies.

2.5 Acknowledgments

We thank the participants in the trial, all COVE study contributors and site investigators, and the members of the data and safety monitoring board (Richard J. Whitley, chair, University of Alabama School of Medicine; Abdel Babiker, MRC Clinical Trials Unit at University College, London; Lisa Angeline Cooper, Johns Hopkins University School of Medicine and Bloomberg School of Public Health; Susan Smith Ellenberg, University of Pennsylvania; Alan Fix, Vaccine Development Global Program Center for Vaccine Innovation and Access PATH; Marie Griffin, Vanderbilt University Medical Center; Steven Joffe, Perelman School of Medicine, Uni- versity of Pennsylvania; Jorge Kalil, Clinics Hospital (HC-FMUSP), Universidade de São Paulo, Brazil; Myron M. Levine, University of Maryland School of Medicine; Malegapuru William Makgoba, University of KwaZulu-Natal; Anastasios A. Tsiatis, North Car- olina State University; and Renee H. Moore, Emory University) and the adjudication committee (Richard J. Hamill, chair, Baylor College of Medicine; Lewis Lipsitz, Harvard Medical School; Eric S. Rosenberg, Massachusetts General Hospital; and Anthony Faugno, Tufts Medical Center) for their critical and timely review of the trial data. We also acknowledge the contribution from the mRNA- 1273 Product Coordination Team from the Biomedical Advanced Research and Development Authority (BARDA) (Robert Bruno, Richard Gorman, Holli Hamilton, Gary Horwith, Chuong Huynh, Nutan Mytle, Corrina Pavetto, Xiaomi Tong, and John Treanor).

Bibliography

[1] Marc Lipsitch. Challenges of vaccine effectiveness and waning studies. *Clin. Infect. Dis.*, 68 (10):1631–1633, May 2019.

[2] G Thomas Ray, Ned Lewis, Nicola P Klein, Matthew F Daley, Marc Lipsitch, and Bruce Fireman. Depletion-of-susceptibles bias in analyses of intra-season waning of influenza vaccine effectiveness. *Clin. Infect. Dis.*, 70(7):1484–1486, March 2020.

[3] Rebecca Kahn, Stephanie J Schrag, Jennifer R Verani, and Marc Lipsitch. Identifying and alleviating bias due to differential depletion of susceptible people in postmarketing evaluations of COVID-19 vaccines. *Am. J. Epidemiol.*, 191(5):800–811, March 2022.

[4] World Health Organization. Evaluation of covid-19 vaccine effectiveness: interim guidance, March 17, 2021. URL `https://www.who.int/publications/i/item/WHO-2019-nCoV-vaccine_effectiveness-measurement-2021.1`. Accessed: 2024-04-13.

[5] World Health Organization. Evaluation of covid-19 vaccine effectiveness in a changing landscape of covid-19 epidemiology and vaccination: Interim guidance, October 3, 2022. URL `https://www.who.int/publications/i/item/WHO-2019-nCoV-vaccine_effectiveness-VE_evaluations-2022.1`. Accessed: 2024-04-13.

[6] Tomáš Fürst, Angelika Bazalová, Tadeáš Fryčák, and Jaroslav Janošek. Does the healthy vaccinee bias rule them all? association of COVID-19 vaccination status and all-cause mortality from an analysis of data from 2.2 million individual health records. *Int. J. Infect. Dis.*, 142 (106976):106976, May 2024.

[7] Kaiser Fung, Mark Jones, and Peter Doshi. Sources of bias in observational studies of covid-19 vaccine effectiveness. *J. Eval. Clin. Pract.*, 30(1):30–36, February 2024.

[8] Korryn Bodner, Jesse Knight, Mackenzie A Hamilton, and Sharmistha Mishra. Testing whether higher contact among the vaccinated can be a mechanism for observed negative vaccine effectiveness. *Am. J. Epidemiol.*, 192(8):1335–1340, August 2023.

[9] William J Hulme, Elizabeth Williamson, Elsie M F Horne, Amelia Green, Helen I McDonald, Alex J Walker, Helen J Curtis, Caroline E Morton, Brian MacKenna, Richard Croker, Amir Mehrkar, Seb Bacon, David Evans, Peter Inglesby, Simon Davy, Krishnan Bhaskaran, Anna Schultze, Christopher T Rentsch, Laurie Tomlinson, Ian J Douglas, Stephen J W Evans, Liam Smeeth, Tom Palmer, Ben Goldacre, Miguel A Hernán, and Jonathan A C Sterne. Challenges in estimating the effectiveness of COVID-19 vaccination using observational data. *Ann. Intern. Med.*, 176(5):685–693, May 2023.

[10] Daniel R Feikin, Melissa M Higdon, Laith J Abu-Raddad, Nick Andrews, Rafael Araos, Yair Goldberg, Michelle J Groome, Amit Huppert, Katherine L O'Brien, Peter G Smith, Annelies Wilder-Smith, Scott Zeger, Maria Deloria Knoll, and Minal K Patel. Duration of effectiveness of vaccines against SARS-CoV-2 infection and COVID-19 disease: results of a systematic review and meta-regression. *Lancet*, 399(10328):924–944, March 2022.

[11] Anastasios A Tsiatis and Marie Davidian. Rejoinder: Estimating vaccine efficacy over time after a randomized study is unblinded. *Biometrics*, 78(3):848–851, September 2022.

[12] Ilari Kuitunen, Mikko Uimonen, Santeri J Seppälä, and Ville T Ponkilainen. COVID-19 vaccination status and testing rates in Finland-A potential cause for bias in observational vaccine effectiveness analysis. *Influenza Other Respi. Viruses*, 16(5):842–845, September 2022.

[13] Paul Glasziou, Kirsten McCaffery, Erin Cvejic, Carys Batcup, Julie Ayre, Kristen Pickles, and Carissa Bonner. Testing behaviour may bias observational studies of vaccine effectiveness. *Off. J. Assoc. Med. Microbiol. Infect. Dis. Can.*, 7(3):242–246, September 2022.

[14] Benjamin Rader, Autumn Gertz, A Danielle Iuliano, Matthew Gilmer, Laura Wronski, Christina M Astley, Kara Sewalk, Tanner J Varrelman, Jon Cohen, Rishika Parikh, Heather E Reese, Carrie Reed, and John S Brownstein. Use of at-home COVID-19 tests - united states, august 23, 2021-march 12, 2022. *MMWR Morb. Mortal. Wkly. Rep.*, 71(13):489–494, April 2022.

[15] Michael L Jackson and Kenneth J Rothman. Effects of imperfect test sensitivity and specificity

on observational studies of influenza vaccine effectiveness. *Vaccine*, 33(11):1313–1316, March 2015.

[16] Manish M Patel, Michael L Jackson, and Jill Ferdinands. Postlicensure evaluation of COVID-19 vaccines. *JAMA*, 324(19):1939–1940, November 2020.

[17] Tom De Smedt, Elizabeth Merrall, Denis Macina, Silvia Perez-Vilar, Nick Andrews, and Kaatje Bollaerts. Bias due to differential and non-differential disease- and exposure misclassification in studies of vaccine effectiveness. *PLoS One*, 13(6):e0199180, June 2018.

[18] Marc Lipsitch, Eric Tchetgen Tchetgen, and Ted Cohen. Negative controls: a tool for detecting confounding and bias in observational studies. *Epidemiology*, 21(3):383–388, May 2010.

[19] Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. A selective review of negative control methods in epidemiology. *Curr. Epidemiol. Rep.*, 7(4):190–202, December 2020.

[20] Joseph A Lewnard, Nathan C Lo, Nimalan Arinaminpathy, Isabel Frost, and Ramanan Laxminarayan. Childhood vaccines and antibiotic use in low- and middle-income countries. *Nature*, 581(7806):94–99, May 2020.

[21] Lola Etievant, Joshua N Sampson, and Mitchell H Gail. Increasing efficiency and reducing bias when assessing HPV vaccination efficacy by using nontargeted HPV strains. *Biometrics*, 79(2):1534–1545, June 2023.

[22] Matt D T Hitchings, Joseph A Lewnard, Natalie E Dean, Albert I Ko, Otavio T Ranzani, Jason R Andrews, and Derek A T Cummings. Use of recently vaccinated individuals to detect bias in test-negative case-control studies of COVID-19 vaccine effectiveness. *Epidemiology*, 33 (4):450–456, July 2022.

[23] Michael L Jackson and Jennifer C Nelson. The test-negative design for estimating influenza vaccine effectiveness. *Vaccine*, 31(17):2165–2168, April 2013.

[24] Natalie E Dean, Joseph W Hogan, and Mireille E Schnitzer. Covid-19 vaccine effectiveness and the test-negative design. *N. Engl. J. Med.*, 385(15):1431–1433, October 2021.

[25] Kristin L Andrejko, Jennifer F Myers, John Openshaw, Nozomi Fukui, Sophia Li, James P Watt, Erin L Murray, Cora Hoover, Joseph A Lewnard, Seema Jain, Jake M Pry, and California COVID-19 Case-Control Study Team. Receipt of COVID-19 and seasonal influenza vaccines in california (USA) during the 2021-2022 influenza season. *Vaccine*, 41(6):1190–1197, February 2023.

[26] Katherine L Anders, Zoe Cutcher, Immo Kleinschmidt, Christl A Donnelly, Neil M Ferguson, Citra Indriani, Peter A Ryan, Scott L O'Neill, Nicholas P Jewell, and Cameron P Simmons. Cluster-randomized test-negative design trials: A novel and efficient method to assess the efficacy of community-level dengue interventions. *Am. J. Epidemiol.*, 187(9):2021–2028, September 2018.

[27] Lindsey R Baden, Hana M El Sahly, Brandon Essink, Dean Follmann, Gregory Hachigian, Cynthia Strout, J Scott Overcash, Susanne Doblecki-Lewis, Jennifer A Whitaker, Evan J Anderson, Kathleen Neuzil, Lawrence Corey, Frances Priddy, Joanne E Tomassini, Mollie Brown, Bethany Girard, Dina Stolman, Veronica Urdaneta, Xiaowei Wang, Weiping Deng, Honghong Zhou, Avika Dixit, Rituparna Das, Jacqueline M Miller, and COVE Trial Consortium. Long-term safety and effectiveness of mRNA-1273 vaccine in adults: COVE trial open-label and booster phases. *Nat. Commun.*, 15(1):7469, August 2024.

[28] Lindsey R Baden, Hana M El Sahly, Brandon Essink, Karen Kotloff, Sharon Frey, Rick Novak, David Diemert, Stephen A Spector, Nadine Rouphael, C Buddy Creech, John McGettigan, Shishir Khetan, Nathan Segall, Joel Solis, Adam Brosz, Carlos Fierro, Howard Schwartz, Kathleen Neuzil, Larry Corey, Peter Gilbert, Holly Janes, Dean Follmann, Mary Marovich, John Mascola, Laura Polakowski, Julie Ledgerwood, Barney S Graham, Hamilton Bennett, Rolando Pajon, Conor Knightly, Brett Leav, Weiping Deng, Honghong Zhou, Shu Han, Melanie Ivarsson, Jacqueline Miller, Tal Zaks, and COVE Study Group. Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine. *N. Engl. J. Med.*, 384(5):403–416, February 2021.

[29] Hana M El Sahly, Lindsey R Baden, Brandon Essink, Susanne Doblecki-Lewis, Judith M Martin, Evan J Anderson, Thomas B Campbell, Jesse Clark, Lisa A Jackson, Carl J Fichtenbaum, Marcus Zervos, Bruce Rankin, Frank Eder, Gregory Feldman, Christina Kennelly, Laurie Han-

Conrad, Michael Levin, Kathleen M Neuzil, Lawrence Corey, Peter Gilbert, Holly Janes, Dean Follmann, Mary Marovich, Laura Polakowski, John R Mascola, Julie E Ledgerwood, Barney S Graham, Allison August, Heather Clouting, Weiping Deng, Shu Han, Brett Leav, Deb Manzo, Rolando Pajon, Florian Schödel, Joanne E Tomassini, Honghong Zhou, Jacqueline Miller, and COVE Study Group. Efficacy of the mRNA-1273 SARS-CoV-2 vaccine at completion of blinded phase. *N. Engl. J. Med.*, 385(19):1774–1785, November 2021.

[30] Xu Shi, Kendrick Qijun Li, and Bhramar Mukherjee. Current challenges with the use of test-negative designs for modeling COVID-19 vaccination and outcomes. *Am. J. Epidemiol.*, 192 (3):328–333, February 2023.

[31] Dean Follmann, Xiaowei Wang, Lindsey R Baden, Hana M El Sahly, Brandon Essink, Peter Gilbert, Holly E Janes, Colleen F Kelley, Megan A Berman, Ian Frank, Eric Chu, Weiping Deng, Frances Priddy, Avika Dixit, Joanne E Tomassini, Rituparna Das, Jacqueline Miller, and Honghong Zhou. Who to boost when: The effect of age and dosing interval on delta and omicron COVID-19 incidence in the open-label phase of the COVE trial. *Open Forum Infect. Dis.*, 11(12):ofae689, December 2024.

[32] Miguel A. Hernán. The hazards of hazard ratios. *Epidemiology*, 21(1):13–15, January 2010.

[33] Torben Martinussen, Stijn Vansteelandt, and Per Kragh Andersen. Subtleties in the interpretation of hazard contrasts. *Lifetime Data Anal.*, 26(4):833–855, October 2020.

[34] Ted Westling, Alex Luedtke, Peter B. Gilbert, and Marco Carone. Inference for Treatment-Specific Survival Curves Using Machine Learning. *Journal of the American Statistical Association*, 0(0):1–13, 2023. ISSN 0162-1459. doi: 10.1080/01621459.2023.2205060. URL `https://doi.org/10.1080/01621459.2023.2205060`.

[35] Peter B Gilbert, David C Montefiori, Adrian B McDermott, Youyi Fong, David Benkeser, Weiping Deng, Honghong Zhou, Christopher R Houchens, Karen Martins, Lakshmi Jayashankar, Flora Castellino, Britta Flach, Bob C Lin, Sarah O'Connell, Charlene McDanal, Amanda Eaton, Marcella Sarzotti-Kelsoe, Yiwen Lu, Chenchen Yu, Bhavesh Borate, Lars W P van der Laan, Nima S Hejazi, Chuong Huynh, Jacqueline Miller, Hana M El Sahly, Lindsey R Baden,

Mira Baron, Luis De La Cruz, Cynthia Gay, Spyros Kalams, Colleen F Kelley, Michele P Andrasik, James G Kublin, Lawrence Corey, Kathleen M Neuzil, Lindsay N Carpp, Rolando Pajon, Dean Follmann, Ruben O Donis, Richard A Koup, Immune Assays Team§, Moderna, Inc. Team§, Coronavirus Vaccine Prevention Network (CoVPN)/Coronavirus Efficacy (COVE) Team§, and United States Government (USG)/CoVPN Biostatistics Team§. Immune correlates analysis of the mRNA-1273 COVID-19 vaccine efficacy clinical trial. *Science*, 375(6576): 43–50, January 2022.

[36] Jonathan Fintzi and Dean Follmann. Assessing vaccine durability in randomized trials following placebo crossover. *Stat. Med.*, 40(27):5983–6007, November 2021.

[37] Paul H C Eilers and Brian D Marx. Flexible smoothing with b-splines and penalties. *Statistical Science*, 11(2):89–121, May 1996.

[38] Clifford M Hurvich, Jeffrey S Simonoff, and Chih-Ling Tsai. Smoothing parameter selection in nonparametric regression using an improved akaike information criterion. *J. R. Stat. Soc. Series B Stat. Methodol.*, 60(2):271–293, July 1998.

[39] Peter B Gilbert, Bryan E Shepherd, and Michael G Hudgens. Sensitivity analysis of per-protocol time-to-event treatment efficacy in randomized clinical trials. *J. Am. Stat. Assoc.*, 108(503):789–800, January 2013.

[40] R L Prentice, J D Kalbfleisch, A V Peterson, Jr, N Flournoy, V T Farewell, and N E Breslow. The analysis of failure times in the presence of competing risks. *Biometrics*, 34(4):541–554, December 1978.

[41] Shruti Khare, Céline Gurry, Lucas Freitas, Mark B Schultz, Gunter Bach, Amadou Diallo, Nancy Akite, Joses Ho, Raphael Tc Lee, Winston Yeo, Gisaid Core Curation Team, and Sebastian Maurer-Stroh. GISAID's role in pandemic response. *China CDC Wkly.*, 3(49): 1049–1051, December 2021.

[42] Ensheng Dong, Hongru Du, and Lauren Gardner. An interactive web-based dashboard to track COVID-19 in real time. *Lancet Infectious Diseases*, 20(5):533–534, May 2020.

[43] Rolando Pajon, Yamuna D Paila, Bethany Girard, Groves Dixon, Katherine Kacena, Lindsey R Baden, Hana M El Sahly, Brandon Essink, Kathleen M Mullane, Ian Frank, Douglas Denhan, Edward Kerwin, Xiaoping Zhao, Baoyu Ding, Weiping Deng, Joanne E Tomassini, Honghong Zhou, Brett Leav, Florian Schödel, and COVE Trial Consortium. Initial analysis of viral dynamics and circulating viral variants during the mRNA-1273 phase 3 COVE trial. *Nat. Med.*, 28(4):823–830, April 2022.

[44] Sonja J Olsen, Amber K Winn, Alicia P Budd, Mila M Prill, John Steel, Claire M Midgley, Krista Kniss, Erin Burns, Thomas Rowe, Angela Foust, Gabriela Jasso, Angiezel Merced-Morales, C Todd Davis, Yunho Jang, Joyce Jones, Peter Daly, Larisa Gubareva, John Barnes, Rebecca Kondor, Wendy Sessions, Catherine Smith, David E Wentworth, Shikha Garg, Fiona P Havers, Alicia M Fry, Aron J Hall, Lynnette Brammer, and Benjamin J Silk. Changes in influenza and other respiratory virus activity during the COVID-19 pandemic - united states, 2020-2021. *MMWR Morb. Mortal. Wkly. Rep.*, 70(29):1013–1019, July 2021.

[45] Gail L Sondermeyer Cooksey, Christina Morales, Lauren Linde, Samuel Schildhauer, Hugo Guevara, Elena Chan, Kathryn Gibb, Jessie Wong, Wen Lin, Brandon J Bonin, Olivia Arizmendi, Tracy Lam-Hine, Ori Tzvieli, Ann McDowell, Kirstie M Kampen, Denise L Lopez, Josh Ennis, Linda S Lewis, Eyal Oren, April Hatada, Blanca Molinar, Matt Frederick, George S Han, Martha Sanchez, Michael A Garcia, Alana McGrath, Nga Q Le, Eric Boyd, Regina M Bertolucci, Jeremy Corrigan, Stephanie Brodine, Michael Austin, William R K Roach, Robert M Levin, Brian M Tyson, Jake M Pry, Kristin J Cummings, Debra A Wadford, and Seema Jain. Severe acute respiratory syndrome coronavirus 2 and respiratory virus sentinel surveillance, california, USA, may 10, 2020-june 12, 2021. *Emerg. Infect. Dis.*, 28(1): 9–19, January 2022.

[46] Heui Man Kim, Eun Jung Lee, Nam-Joo Lee, Sang Hee Woo, Jeong-Min Kim, Jee Eun Rhee, and Eun-Jin Kim. Impact of coronavirus disease 2019 on respiratory surveillance and explanation of high detection rate of human rhinovirus during the pandemic in the republic of korea. *Influenza Other Respi. Viruses*, 15(6):721–731, November 2021.

[47] Sangshin Park, Ian C Michelow, and Young June Choe. Shifting patterns of respiratory virus

activity following social distancing measures for coronavirus disease 2019 in south korea. *J. Infect. Dis.*, 224(11):1900–1906, December 2021.

[48] Asmaa El-Heneidy, Robert S Ware, Jennifer M Robson, Sarah G Cherian, Stephen B Lambert, and Keith Grimwood. Respiratory virus detection during the COVID-19 pandemic in queensland, australia. *Aust. N. Z. J. Public Health*, 46(1):10–15, February 2022.

[49] Ilari Kuitunen, Miia Artama, Marjut Haapanen, and Marjo Renko. Rhinovirus spread in children during the COVID-19 pandemic despite social restrictions-a nationwide register study in finland. *J. Med. Virol.*, 93(10):6063–6067, October 2021.

[50] Ana Karolina Antunes Eisen, Juliana Schons Gularte, Meriane Demoliner, Vyctoria Malayhka de Abreu Goés Pereira, Fágner Henrique Heldt, Micheli Filippi, Paula Rodrigues de Almeida, Alana Witt Hansen, Juliane Deise Fleck, and Fernando Rosado Spilki. Low circulation of influenza a and coinfection with SARS-CoV-2 among other respiratory viruses during the COVID-19 pandemic in a region of southern brazil. *J. Med. Virol.*, 93(7):4392–4398, July 2021.

[51] Kristal An Agrupis, Annavi Marie G Villanueva, Ana Ria Sayo, Jezreel Lazaro, Su Myat Han, Alyannah C Celis, Shuichi Suzuki, Ann Celestyn Uichanco, Jocelyn Sagurit, Rontgene Solante, Lay-Myint Yoshida, Koya Ariyoshi, and Chris Smith. If not COVID-19 what is it? analysis of COVID-19 versus common respiratory viruses among symptomatic health care workers in a tertiary infectious disease referral hospital in manila, philippines. *Trop. Med. Infect. Dis.*, 6 (1):39, March 2021.

[52] Eric J Chow, Timothy M Uyeki, and Helen Y Chu. The effects of the COVID-19 pandemic on community respiratory virus activity. *Nat. Rev. Microbiol.*, 21(3):195–210, March 2023.

[53] Thomas Hills, Nethmi Kearns, Ciléin Kearns, and Richard Beasley. Influenza control during the COVID-19 pandemic. *Lancet*, 396(10263):1633–1634, November 2020.

[54] Luzhao Feng, Ting Zhang, Qing Wang, Yiran Xie, Zhibin Peng, Jiandong Zheng, Ying Qin, Muli Zhang, Shengjie Lai, Dayan Wang, Zijian Feng, Zhongjie Li, and George F Gao. Impact

of COVID-19 outbreaks and interventions on influenza in china and the united states. *Nat. Commun.*, 12(1):3249, May 2021.

[55] Sonja J Olsen, Eduardo Azziz-Baumgartner, Alicia P Budd, Lynnette Brammer, Sheena Sullivan, Rodrigo Fasce Pineda, Cheryl Cohen, and Alicia M Fry. Decreased influenza activity during the COVID-19 pandemic - united states, australia, chile, and south africa, 2020. *MMWR Morb. Mortal. Wkly. Rep.*, 69(37):1305–1309, September 2020.

[56] Karen L Laurie and Steve Rockman. Which influenza viruses will emerge following the SARS-CoV-2 pandemic? *Influenza Other Respi. Viruses*, 15(5):573–576, September 2021.

[57] Fernanda Hammes Varela, Marcelo Comerlato Scotta, Márcia Polese-Bonatto, Ivaine Tais Sauthier Sartor, Charles Francisco Ferreira, Ingrid Rodrigues Fernandes, Gabriela Oliveira Zavaglia, Walquiria Aparecida Ferreira de Almeida, Denise Arakaki-Sanchez, Leonardo Araújo Pinto, Gisele Alsina Nader Bastos, Luiz Antônio Nasi, Maicon Falavigna, Paulo Márcio Pitrez, Renato T Stein, and COVIDa study group. Absence of detection of RSV and influenza during the COVID-19 pandemic in a brazilian cohort: Likely role of lower transmission in the community. *J. Glob. Health*, 11(05007):05007, March 2021.

[58] Stefano Tempia, Sibongile Walaza, Jinal N Bhiman, Meredith L McMorrow, Jocelyn Moyes, Thulisa Mkhencele, Susan Meiring, Vanessa Quan, Kate Bishop, Johanna M McAnerney, Anne von Gottberg, Nicole Wolter, Mignon Du Plessis, Florette K Treurnicht, Orienka Hellferscee, Halima Dawood, Fathima Naby, Ebrahim Variava, Comfort Siwele, Neydis Baute, Jeremy Nel, Gary Reubenson, Heather J Zar, and Cheryl Cohen. Decline of influenza and respiratory syncytial virus detection in facility-based surveillance during the COVID-19 pandemic, south africa, january to october 2020. *Euro Surveill.*, 26(29), July 2021.

[59] Jean-Sebastien Casalegno, Dominique Ploin, Aymeric Cantais, Elsa Masson, Emilie Bard, Martine Valette, Remi Fanget, Sandrine Couray Targe, Anne-Florence Myar-Dury, Muriel Doret-Dion, Mona Massoud, Gregory Queromes, Philippe Vanhems, Olivier Claris, Marine Butin, Sylvie Pillet, Florence Ader, Sylvie Bin, Alexandre Gaymard, Bruno Lina, Florence Morfin, VRS study group in Lyon, Etienne Javouhey, and Yves Gillet. Characteristics of the

delayed respiratory syncytial virus epidemic, 2020/2021, rhône loire, france. *Euro Surveill.*, 26 (29), July 2021.

[60] Daniel K Yeoh, David A Foley, Cara A Minney-Smith, Andrew C Martin, Ariel O Mace, Chisha T Sikazwe, Huong Le, Avram Levy, Christopher C Blyth, and Hannah C Moore. Impact of coronavirus disease 2019 public health measures on detections of influenza and respiratory syncytial virus in children during the 2020 australian winter. *Clin. Infect. Dis.*, 72 (12):2199–2202, June 2021.

[61] Emi Takashita, Chiharu Kawakami, Tomoko Momoki, Miwako Saikusa, Kouhei Shimizu, Hiroki Ozawa, Makoto Kumazaki, Shuzo Usuku, Nobuko Tanaka, Ichiro Okubo, Hiroko Morita, Shiho Nagata, Shinji Watanabe, Hideki Hasegawa, and Yoshihiro Kawaoka. Increased risk of rhinovirus infection in children during the coronavirus disease-19 pandemic. *Influenza Other Respi. Viruses*, 15(4):488–494, July 2021.

[62] Sujin Yum, Kwan Hong, Sangho Sohn, Jeehyun Kim, and Byung Chul Chun. Trends in viral respiratory infections during COVID-19 pandemic, south korea. *Emerg. Infect. Dis.*, 27(6): 1685–1688, 2021.

[63] Kevin Zhang, Avika Misra, Patrick J Kim, Seyed M Moghadas, Joanne M Langley, and Marek Smieja. Rapid disappearance of influenza following the implementation of COVID-19 mitigation measures in hamilton, ontario. *Can. Commun. Dis. Rep.*, 47(4):202–209, May 2021.

[64] Seattle Flu Alliance. Pathogens. `https://seattleflu.org/pathogens`, 2026. Accessed: 2026-05-07.

[65] Centers for Disease Control and Prevention. National respiratory and enteric virus surveillance system (nrevss). `https://www.cdc.gov/nrevss/php/dashboard/index.html#cdc_tools_technologies_download-viruses-monitored-in-nrevss`, 2024. Published online 2024. Accessed May 1, 2025.

[66] Cary Wu and Yue Qian. The gender peak effect: Women are most vulnerable to infections during COVID-19 peaks. *Front. Public Health*, 10:937179, August 2022.

[67] Holly Janes, Leigh H Fisher, Jia Jin Kee, Lalitha Parameswaran, Paul A Goepfert, Ann R Falsey, James Ludwig, Craig A Magaret, Peter B Gilbert, James G Kublin, Nadine Rouphael, Magdalena E Sobieszczyk, Hana M El Sahly, Lindsey R Baden, Beatriz Grinsztejn, Stephen R Walsh, Glenda E Gray, Karen L Kotloff, Cynthia L Gay, Alexander L Greninger, Milagritos D Tapia, E Adrianne Hammershaimb, Frances H Priddy, Justin A Green, Frank Struyf, Lisa Dunkle, Kathleen M Neuzil, Lawrence Corey, and Yunda Huang. Association between SARS-CoV-2 viral load and COVID-19 vaccination in 4 phase 3 trials. *J. Infect. Dis.*, 230(6):1384–1389, December 2024.

[68] Margaret K Doll, Stacy M Pettigrew, Julia Ma, and Aman Verma. Effects of confounding bias in coronavirus disease 2019 (COVID-19) and influenza vaccine effectiveness test-negative designs due to correlated influenza and COVID-19 vaccination behaviors. *Clin. Infect. Dis.*, 75(1):e564–e571, August 2022.

[69] Sheena G Sullivan, Eric J Tchetgen Tchetgen, and Benjamin J Cowling. Theoretical basis of the test-negative study design for assessment of influenza vaccine effectiveness. *Am. J. Epidemiol.*, 184(5):345–353, September 2016.

[70] Leah I B Andrews, M Elizabeth Halloran, Kathleen M Neuzil, Lars van der Laan, Yunda Huang, Jessica Andriesen, Mayur Patel, Leigh H Fisher, Holly Janes, Nadine Rouphael, Stephen R Walsh, Deborah A Theodore, Hong-Van Tieu, Magdalena Sobieszczyk, Hana M El Sahly, Lindsey R Baden, Ann R Falsey, Thomas B Campbell, Colleen F Kelley, Catherine Mary Healy, Lilly Immergluck, Benjamin Luft, Ian Hirsch, Guy de Bruyn, Carla Truyers, Frances Priddy, Kelsey M Sumner, Brendan Flannery, Dean Follmann, Peter B Gilbert, and COVID-19 Prevention Network (CoVPN). Evaluating the test-negative design for COVID-19 vaccine effectiveness using randomized trial data: A secondary cross-protocol analysis of 5 randomized clinical trials. *JAMA Network Open*, 8(5):e2512763, May 2025.

[71] Benjamin F Arnold, Ayse Ercumen, Jade Benjamin-Chung, and John M Colford, Jr. Brief report: Negative controls to detect selection bias and measurement bias in epidemiologic studies. *Epidemiology*, 27(5):637–641, September 2016.

[72] M. Lipsitch, E. Goldstein, G.T. Ray, and B. Fireman. Depletion-of-susceptibles bias in influenza vaccine waning studies: how to ensure robust results. *Epidemiology & Infection*, 2019.

[73] Jocelyne Piret and Guy Boivin. Viral interference between respiratory viruses. *Emerg. Infect. Dis.*, 28(2):273–281, February 2022.

[74] Centers for Medicare & Medicaid Services. Fy 2025 hospital inpatient prospective payment system (ipps) and long-term care hospital prospective payment system (ltch pps) final rule–cms-1808-f. `https://www.cms.gov/newsroom/fact-sheets/fy-2025-hospital-inpatient-prospective-payment-system-ipps-and-long-term-care-hospital-pr` 2024. Published online August 1, 2024. Accessed February 1, 2025.

[75] Chan Park, David B Richardson, and Eric J Tchetgen Tchetgen. Single proxy control. *Biometrics*, 80(2):ujae027, March 2024.

[76] David B Richardson, Dominique Laurier, Mary K Schubauer-Berigan, Eric Tchetgen Tchetgen, and Stephen R Cole. Assessment and indirect adjustment for confounding by smoking in cohort studies using relative hazards models. *Am. J. Epidemiol.*, 180(9):933–940, November 2014.

[77] Eric J Tchetgen Tchetgen, Chan Park, and David B Richardson. Universal difference-in-differences for causal inference in epidemiology. *Epidemiology*, 35(1):16–22, January 2024.

[78] Kendrick Qijun Li, Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. Double negative control inference in test-negative design studies of vaccine effectiveness. *J. Am. Stat. Assoc.*, 119(547): 1859–1870, 2024.

[79] Yifan Cui, Hongming Pu, Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. Semiparametric proximal causal inference. *J. Am. Stat. Assoc.*, 119(546):1348–1359, April 2024.

[80] Paul N Zivich, Stephen R Cole, Jessie K Edwards, Grace E Mulholland, Bonnie E Shook-Sa, and Eric J Tchetgen Tchetgen. Introducing proximal causal inference for epidemiologists. *Am. J. Epidemiol.*, 192(7):1224–1227, July 2023.

[81] Peter B Gilbert, Youyi Fong, Avi Kenny, and Marco Carone. A controlled effects approach to assessing immune correlates of protection. *Biostatistics*, 24(4):850–865, October 2023.

Chapter 3

DEBIASING HAZARD-BASED, TIME-VARYING VACCINE EFFECTIVENESS ESTIMATES USING NEGATIVE CONTROL INFECTIONS

Summary: *Understanding how vaccine effectiveness (VE) changes over time can provide evidence-based guidance for public health decision making. While commonly reported by practitioners, time-varying VE estimates obtained using Cox regression are vulnerable to hidden biases. To address these limitations, we leverage vaccine-irrelevant, negative control infections to identify hazard-based, time-varying VE under unmeasured confounding and selection bias. We articulate assumptions under which our approach identifies a causal effect of an intervention deferring vaccination and interaction with the community in which infections circulate. We develop sieve and efficient influence curve-based estimators and discuss imposing monotone shape constraints and estimating VE against multiple variants. As a case study, we examine the observational booster phase of the Coronavirus Vaccine Efficacy (COVE) trial of the Moderna mRNA-1273 COVID-19 vaccine which used symptom-triggered multiplex PCR testing to identify acute respiratory illnesses (ARIs) caused by SARS-CoV-2 and 20 off-target pathogens previously identified as compelling negative controls for COVID-19. Accounting for vaccine-irrelevant ARIs supported that the mRNA-1273 booster was more effective and durable against Omicron COVID-19 than suggested by naïve Cox regression. Our work offers a way to mitigate bias in hazard-based, time-varying treatment effects in randomized and non-randomized studies using negative controls.*

3.1 Introduction

A key goal of vaccine research is to develop novel vaccines or adapt existing vaccines to be long-lasting and broadly protective against many pathogen strains. Understanding the breadth and durability of vaccine effectiveness (VE) can inform whether and when to offer vaccines to the public. For example, evidence of waning influenza vaccine effectiveness supports the practice of vaccinating high-risk groups close to the start of flu season rather than well in advance.[1,2] Hence, *time-varying VE*, or vaccine effectiveness in preventing disease X-days after vaccination, is of interest to both

vaccine science and public health. Many practitioners estimate time-varying VE by fitting a Cox proportional hazards model with a coefficient that varies in time-since-vaccination.[3–7] However, to draw valid conclusions about VE, the vaccinated and unvaccinated populations being compared should be otherwise *exchangeable*, meaning they have the same distribution of risk factors. As we explain below, time-varying VE estimates based on hazard functions can be misleading, because they are vulnerable to violations of exchangeability resulting from both structural features of the hazard function and aspects of study design.

Even in randomized trials, time-varying hazard ratios can provide misleading evidence of VE waning due to a phenomenon referred to as "depletion of susceptibles bias" among epidemiologists (see Figure 1.1 for an illustration and Section 1.2 for extended discussion).[8,9] Even though randomization guarantees balanced risk factors between vaccinated and placebo cohorts at baseline, when a vaccine is protective, "high-risk" placebo recipients will succumb to infection at a faster rate than "high-risk" vaccinated participants. Since the hazard function conditions on event-free survival, the time-varying hazard ratio can become distorted by time-varying imbalances in infection risk between vaccinated and unvaccinated disease-free cohorts, which can lead to exaggerated magnitude of waning VE. Using a potential outcomes framework, Hernán[10] clarified that depletion of susceptibles bias is a structural limitation of the hazard function itself, which he termed a "built-in selection bias." In an effort to mitigate depletion of susceptibles, analysts typically adjust for known important risk factors, for example by using serological testing to account for natural immunity from prior infections[11]. Others have proposed alternative vaccine trial designs[9,12] and causal estimands[13] that avoid differential depletion of susceptibles by exploiting natural or hypothetical interventions that delay exposure to the target pathogen until a future date. While these approaches are promising, they require strong assumptions that are unlikely to be satisfied in many contexts.

Moreover, conducting long, randomized vaccine trials to evaluate waning VE may not be feasible due to resource limitations, lack of clinical equipoise, or pathogen evolution in the population[14]. Alternatively, observational studies contain rich data on how vaccine effectiveness varies across different settings, variants, and over time. However, VE estimates from observational studies are particularly vulnerable to bias. Myriad sources of confounding can impact observational vaccine studies[15,16] including differences in behaviors (such as mask-wearing and social distancing) influenc-

ing contact rates with infected community members[17,18], immune function influencing per-contact susceptibility to infection[11,17], and healthcare seeking behaviors impacting case ascertainment[19,20]. These concerns are not theoretical: several observational COVID-19 vaccine studies in winter 2021 reported low and negative VE estimates, which are hypothesized to be the result of hidden bias.[21]

Common strategies for bias mitigation like observed confounder adjustment are often ill-suited to vaccine studies because the relevant sources of bias are hard to measure, overlapping, and may vary over time. For example, a key confounder is an individual's "exposure", often conceptualized as the rate of contact with infected individuals in the community at time t.[17] While exposure can be described conceptually, it is difficult to measure and incorporate into statistical models. For community-transmitted respiratory pathogens, the only requirement for transmission is spatiotemporal proximity, so there is no well-defined, directly observable event that corresponds to "contact with an infected individual".[22–24] Additionally, a participant's exposure can vary over time due to a constellation of factors including changes in behaviors, travel, and non-pharmaceutical intervention adherence. In practice, analysts therefore adjust for measured variables such as demographic and occupational characteristics (e.g., age or healthcare worker status), with the hope that they are correlated with the underlying, unmeasured exposure process.

In this chapter, we will illustrate the potential of using vaccine-irrelevant, negative control infection endpoints to more reliably infer time-varying vaccine effectiveness in the presence of unmeasured sources of selection and confounding bias. As a case study, we will focus on the booster phase of the Coronavirus Vaccine Efficacy (COVE) study, a phase 3, randomized, double-blind, placebo-controlled study of the Moderna mRNA-1273 COVID-19 vaccine in the US adults aged $\geq$18 years. Details on the COVE study can be found in Section 2.2.1 and in prior publications.[25–27] On September 23, 2021, the study protocol was amended to offer a single 50-μg mRNA-1273 booster to participants. COVE's "booster phase" spanned from September 23, 2021 to the data cutoff date (April 5, 2022) and followed participants for a median of 5.3 months post-boost. The dominant SARS-CoV-2 lineage was Delta followed by Omicron, and COVID-19 trends in COVE mirrored those observed in the United States (US) (see Figure 2.2).

3.2 Methods

3.2.1 Notation

Consider a study that measures two infections: a vaccine-preventable infection (denoted $J = 1$) and a vaccine-irrelevant infection (denoted $J = 0$). Because infection rates vary over calendar time, all subsequent developments use a calendar time scale. Let $t = 0$ refer to the date of study start. Throughout we will assume non-informative censoring of infection times. Define T as the earliest infection date of either type and V as the vaccination date. Let $\mathcal{U}(t) := \{U(s) : s \leq t\}$ denote a latent frailty process summarizing the history of a potentially large, time-varying collection of participant characteristics that influence infection risk, although we do not need to specify what they are (see next subsection for additional discussion). We allow the process to evolve arbitrarily and assume instantaneous hazard of infection at time t depends on the frailty process history only through its current value $U(t)$. We define the cause-specific hazard for infection type $J = j$ as

$$h_j(t \mid V = v, U(t) = u) := \lim_{\epsilon \downarrow 0} \frac{P(J = j, T \in [t, t + \epsilon) \mid V = v, U(t) = u, T \geq t)}{\epsilon}$$

where we omit stratification or adjustment for measured covariates for brevity, although they can be readily incorporated. In the developments below, let $E_{U(t) \sim F_{U(t)}}[\cdot]$ denote the expectation with respect to $U(t)$ drawn from a distribution $F_{U(t)}$.

3.2.2 Time-varying, frailty-standardized VE

Our parameter of interest is the marginal *time-varying, frailty-standardized VE curve*, defined as one minus the hazard ratio at time t, standardized to the frailty distribution of vaccinated survivors, denoted by $F^1_{U(t)}$.

$$\mathrm{VE}_1(t, v) := 1 - \frac{E_{U(t) \sim F^1_{U(t)}}[h_1(t \mid V = v, U(t))]}{E_{U(t) \sim F^1_{U(t)}}[h_1(t \mid V > t, U(t))]} \tag{3.1}$$

While the canonical time-varying hazard ratio suffers from bias and unclear interpretation, 3.1 is interpretable as a hazard-based treatment effect parameter defined among the cohort of vaccinated participants who were disease-free at time t.

How does one interpret the latent frailty process? $U(t)$ can be conceptualized as a prognostic summary of a potentially high-dimensional, time-varying collection of variables that influence infec-

tion risk – both unmeasured confounders and precision variables. For example, under the conceptualization of infection risk of Halloran et al.[28], $U(t)$ represents the product of (i) the participant-level contact rate with infected individuals in the community, (ii) the per-contact probability of infection, and (iii) the per-infection case ascertainment probability at time t.[18] Each component may be influenced by a complex array of behavioral and immunological characteristics, but their combined effect on infection risk can be summarized through $U(t)$. By marginalizing the numerator and denominator with respect to a common distribution of risk factors encoded by $U(t)$, the frailty-standardized VE avoids confounding and selection bias.

3.2.3 Canonical model: a multiplicative, time-varying frailty model

To provide a concrete model, we extend the hazard parametrizations of Fintzi and Follmann[4] to include the time-varying unmeasured frailty $U(t)$ which affects the hazard on a multiplicative scale.[24] Define $h_{01}(t)$ and $h_{00}(t)$ as the cause-specific baseline hazards of vaccine-preventable and irrelevant infections respectively.

$$\begin{aligned} h_1(t \mid V, U(t)) &= h_{01}(t)U(t)\exp\{I(V \leq t)f(t-V)\}, \\ h_0(t \mid V, U(t)) &= h_{00}(t)U(t). \end{aligned} \tag{3.2}$$

Where the function $f(\cdot)$ encodes the time-varying vaccine effectiveness on the log hazard ratio scale. This model is more flexible than static frailty models, as they allow unmeasured, time-varying factors to influence infection risk. The frailty $U(t)$ can introduce selection bias in naïve hazard ratios when $U(t)$ is omitted from the model. Moreover, different associations between $\mathcal{U}(t)$ and V can reflect different forms of confounding. For example, high-risk persons may receive vaccination earlier due to personal knowledge of risk or prioritization by mass-vaccination campaigns; this is captured by a negative association between V and $U(t)$. Another example is behavioral disinhibition, where participants cease non-pharmaceutical interventions (e.g., mask wearing, social distancing) after vaccination, thereby increasing their risk of infection. This is captured by an association between $\mathcal{U}(t)$ and V, where values of $U(t)$ tend to be larger for the post-vaccination period ($t > V$).

Below, we will use this model to ground our framework and examine requisite identification assumptions.

3.2.4 Identification of time-varying, frailty-standardized VE

In all subsequent developments, we assume the hazards for each infection type are nonzero such that the associated hazard ratios are well-defined. This requires that both vaccine preventable and irrelevant infections co-circulate during a common calendar time period. Suppose we make the following assumptions.

Assumption 5 (Weak overlap)**.** *For all $t > 0$ and all u in the support of $F^1_{U(t)}$, $P_{0,V}(t < V \mid U(t) = u) > 0$.*

Assumption 6 (Negative control outcome)**.** *For any $t > 0$, assume $h_0(t \mid V, U(t)) = h_0(t \mid U(t))$ with probability 1.*

Assumption 7 (VE varies in time-since-vaccination)**.** *For all $t > 0, v > 0$ with positive support,*

$$\mathrm{VE}_1(t, v) = \mathrm{VE}_1(t - v).$$

Assumption 8 (Bias equivalence)**.** *Define $F^1_{U(t)}$ and $F^0_{U(t)}$ as the distributions of $U(t)$ among disease-free vaccinated participants $(T \geq t, V = v)$ and disease-free unvaccinated participants $(T \geq t, V > t)$ respectively. For all $t \geq 0$, assume*

$$\frac{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V > t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_1(t \mid V > t, U(t))]} = \frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V > t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_0(t \mid V > t, U(t))]}$$

Assumption 5 requires sufficient variability in vaccination statuses within levels of the frailty $U(t)$. This assumption ensures identifiability and that all conditional expectations and hazards are well-defined. Assumption 5 could be violated when all study participants are rapidly vaccinated. In Appendix B.8.2, we discuss ways to recover identifiability of model parameters using external population-level disease surveillance data under no/low overlap in vaccination statuses.

Assumption 6 asserts that the hazard of vaccine-irrelevant infection is independent of vaccination date after conditioning on susceptibility. The high antigenic specificity of the adaptive immune response and the absence of a vaccine effect on off-target ARIs in COVE's blinded phase,[29] provide support for this assumption. While Assumption 6 is biologically credible, in some cases, a small

vaccine effect on irrelevant infections cannot be ruled out. In Appendix B.8.2, we introduce a sensitivity analysis to examine the robustness of our results to potential vaccine effectiveness against the irrelevant infection. Assumption 6 is satisfied in our canonical model 3.2, as there is no term pertaining to vaccination date V in the hazard model for the irrelevant infections.

Assumption 7 asserts that VE varies in time-since-vaccination, and precludes VE variation in calendar time (either vaccination date or the current date). A canonical situation where Assumption 7 could be violated and VE could depend on calendar time if a new viral variant emerges for which the vaccine is less effective.[30] To address this limitation in our case study, we conduct analyses stratified by time periods where Delta or Omicron were predominant, and adapt our approach to estimate variant-specific, time-varying VE against Delta and Omicron COVID-19 (see Section 3.2.6 and Appendix B.8.3 for details). In the canonical model 3.2, Assumption 7 is satisfied because the vaccine effectiveness is entirely parametrized through the function $f(\cdot)$, which varies exclusively in time-since-vaccination.

Assumption 8 is interpreted as follows: among individuals remaining event-free at time t, differences in the latent frailty distributions between vaccinated and unvaccinated groups (due to confounding and selection) impact the vaccine-preventable and irrelevant hazards equivalently on the hazard ratio scale. The dependence of vaccine-preventable and vaccine-irrelevant infections on a common unmeasured causes is a standard assumption referred to as *U-comparability.*[31,32] Assumption 8 is akin to equiconfounding assumption on the hazard ratio scale[33–35] but also accounts for selection bias caused by differential depletion of susceptibles. Assumption 8 will be most plausible when the vaccine-preventable and vaccine-irrelevant infections depend on a common set of unmeasured factors that exert similar effects on the hazards of each endpoint. In Appendix B.1.2, we introduce a sensitivity analysis where the bias inflicted on the vaccine-preventable and vaccine-irrelevant outcomes differ by a multiplicative factor. Focusing on the motivating example 3.2 and applying Tower Law confirms that Assumption 8 is satisfied, as both the vaccine-preventable and vaccine-irrelevant hazard ratios equal the ratio of mean frailties among vaccinated survivors and unvaccinated survivors.

$$\frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V > t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_0(t \mid V > t, U(t))]} = \frac{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V > t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_1(t \mid V > t, U(t))]} = \frac{E_{U(t)\sim F^1_{U(t)}}[U(t)]}{E_{U(t)\sim F^0_{U(t)}}[U(t)]}$$

As a starting point for identification, we examine the marginal time-varying hazard ratio for vaccine-preventable infections. By assumption 7 and arguments in Appendix B.1.1, we can express the marginal time-varying hazard ratio as follows.

$$\frac{h_1(t \mid V = v)}{h_1(t \mid V > t)} = \{1 - \mathrm{VE}_1(t - v)\} \times \frac{E_{U(t) \sim F^1_{U(t)}}[h_1(t \mid V > t, U(t))]}{E_{U(t) \sim F^0_{U(t)}}[h_1(t \mid V > t, U(t))]}$$

Hence, the vaccine-preventable hazard ratio equals the target parameter times a bias term resulting from marginalizing over frailty distributions $F^1_{U(t)} \overset{d}{\neq} F^0_{U(t)}$ that differ due to differential depletion of susceptibles or confounding in the case of an observational study.

While the frailty distribution shift is not directly observed, under assumption 6, the cause-specific hazard ratio for vaccine-irrelevant infections at time t indirectly reveals the frailty imbalance.

$$\frac{h_0(t \mid V = v)}{h_0(t \mid V > t)} = \frac{E_{U(t) \sim F^1_{U(t)}}[h_0(t \mid V > t, U(t))]}{E_{U(t) \sim F^0_{U(t)}}[h_0(t \mid V > t, U(t))]}$$

If the effect of frailty imbalance is symmetric on both vaccine-preventable and vaccine-irrelevant infection endpoints (assumption 8), then the frailty-standardized, time-varying VE is identified as a (one minus) the ratio of cause-specific hazard ratios.

$$\mathrm{VE}_1(t - v) = 1 - \left(\frac{h_1(t \mid V = v)/h_0(t \mid V = v)}{h_1(t \mid V > t)/h_0(t \mid V > t)} \right)$$

In principle, one could estimate the time-varying, frailty-standardized VE by estimating each cause-specific hazard or hazard ratio nonparametrically (e.g., using kernel smoothing) and compute the contrast.[36] However, this approach will suffer from slow convergence rates and instability in finite samples. Using conditional probability arguments (Appendix B.1.1), we can re-express the ratio of cause-specific hazard ratios as an odds ratio of infection causes conditional on date-t infection.

$$\mathrm{VE}_1(t - v) = 1 - \left(\frac{P(J = 1 \mid T = t, V = v)/P(J = 0 \mid T = t, V = v)}{P(J = 1 \mid T = t, V > t)/P(J = 0 \mid T = t, V > t)} \right) \tag{3.3}$$

In Appendix B.2, we examine partial and profile likelihoods for estimating time-varying VE under the multiplicative frailty model in our motivating example 3.2. Interestingly, we show that the time-varying VE parameter identified under partial and profile likelihoods in the multiplicative frailty model is exactly equivalent to the more general odds-ratio identification result which does not rely

on particular hazard parametrizations. Hence, our identification result is "assumption lean," as it can be readily interpreted using partial/profile likelihood if the multiplicative frailty model is correctly specified, but still retain interpretability when the more general assumptions hold.

3.2.5 Inference

Since the estimand is identified by a contrast of odds of infection type among individuals who experience infection at time t, it is natural to model the conditional probability of the infection type indicator J given vaccination status V and infection time t. This leads to an equivalent expression of the target parameter using a logistic regression model:

$$\log\left\{\frac{P(J=1 \mid V,T=t)}{P(J=0 \mid V,T=t)}\right\} = I(V \leq t)f(t-V) + \alpha_0(t) \tag{3.4}$$

Where $f(t-v) = \log\{1-\mathrm{VE}_1(t-v)\}$ parametrizes the time-varying, frailty-standardized VE and $\alpha_0(t) := \log\{P(J=1 \mid V>t, T=t)/P(J=0 \mid V>t, T=t)\}$ is the log odds of vaccine-preventable infection among those unvaccinated and infected at time t, which captures the background composition of infection types among the unvaccinated cases. We consider $f(\cdot)$ as the parameter of primary interest, while $\alpha_0(t)$ is a nuisance parameter.

One strategy for inference in 3.4 is to propose parametric models for $f(\cdot)$ and $\alpha_0(t)$ and maximize the associated logistic likelihood. However, parametric models are likely too simplistic in epidemic settings where $\alpha_0(t)$ fluctuates over calendar time. An alternative is to model both $f(\cdot)$ and $\alpha_0(t)$ nonparametrically. However, purely nonparametric modeling leads to undesirable downstream statistical properties: $f(\cdot)$ is not a pathwise differentiable parameter in 3.4, meaning that root-n consistent and asymptotically normal (CAN) estimators are not generally attainable. Moreover, nonparametric modeling can lead to noisy estimates in finite samples that ignore scientific knowledge about the shape of VE function.

We propose estimation in a working *semiparametric* logistic regression (SLR) model[37–40] which recovers pathwise differentiability and root-n CAN estimators of time-varying, frailty-standardized VE.

$$\log\left\{\frac{P(J=1|V,T=t)}{P(J=0|V,T=t)}\right\} = I(V \leq t)\beta_0^T \boldsymbol{\psi}(t-V) + \alpha_0(t) \tag{3.5}$$

The SLR assumes that $f(\cdot)$ can be spanned by a finite-dimensional (i.e., parametric) basis: $f(\cdot) =$

$\beta_0^T \boldsymbol{\psi}$ for $\beta_0 \in \mathbb{R}^d$ and collection of d basis functions $\boldsymbol{\psi} = \{\psi_1, \ldots, \psi_d\}$ with support on $\mathbb{R}^+$. However, we leave $\alpha_0(t)$ as an unspecified function of calendar time, accommodating complicated fluctuations in the relative background abundances of vaccine-preventable and irrelevant infections. Then, we estimate the frailty-standardized, time-varying VE as follows.

$$\widehat{\mathrm{VE}}_1(t-v) = 1 - \exp\{I(v \leq t)\hat{\beta}^T \boldsymbol{\psi}(t-v)\} \tag{3.6}$$

Below we introduce two strategies for inference in the SLR model shown in 3.5.

Inference via method of sieves

One strategy for inference in 3.5 is to approximate the unknown function $\alpha_0(t)$ with a basis that grows with the sample size, an approach known as sieve estimation.[41] Suppose we approximate $\alpha_0(t)$ via

$$\alpha_0(t) \approx \sum_{j=1}^{M} \alpha_j \phi_{j,v}(t)$$

where $\{\phi_{j,K,v}\}_{j=1}^{M_n}$ is a collection of B-splines of degree v with K_n equally placed knots and coefficients $\alpha_{M_n}^T \in \mathbb{R}^{M_n}$. Let $\hat{\beta}_{K,n}$ refer to the maximum likelihood estimator for β_0 when $\alpha_0(t)$ is approximated using the sieve of dimension M_n. As the sample size grows, we permit the number of knots $K_n \to \infty$, enabling increasingly accurate approximation of the unknown function. If $\alpha_0(t)$ is sufficiently smooth, the sieve's approximation error will vanish at a fast enough rate such that $\hat{\beta}_{K,n}$ will be CAN for β_0. Formally, we require $\alpha_0(t)$ to be uniformly bounded and belong to a Hölder class $\mathcal{G}(v_0, \gamma, c_0)$, consisting of functions g satisfying

$$|g^{(v_0)}(x_1) - g^{(v_0)}(x_0)| \leq c_0 |x_1 - x_0|^\gamma$$

For arbitrary finite γ and c_0 and for all $x_1, x_0 \in [0, 1]$. Some common examples of Hölder classes are Lipschitz functions ($v_0 = 0, \gamma = 1$), functions with bounded derivatives ($v_0 = 1, \gamma = 0$), and functions with bounded higher order derivatives ($v_0 = r, \gamma = 0$). Let $p = v_0 + \gamma$ refer to the *smoothness index* of the function class and let $M_n \asymp n^\lambda$ where λ is the *sieve dimension growth rate*.

Before presenting asymptotic results, we introduce some notation. Define the conditional probability of vaccine-preventable infection on a given vaccination date as follows.

$$Q_0(v,t) = \mathbb{P}(J=1|V=v,T=t) = \frac{\exp\{I(v\le t)\beta_0^T\boldsymbol{\psi}(t-v)+\alpha_0(t)\}}{1+\exp\{I(v\le t)\beta_0^T\boldsymbol{\psi}(t-v)+\alpha_0(t)\}}$$

Define the following d-dimensional nuisance parameter vector which models the vaccination date.

$$\begin{aligned}\boldsymbol{r}_0(t) &:= \frac{\mathbb{E}_V[I(V\le T)\boldsymbol{\psi}(T-V)Q_0(1-Q_0)(V,T)|T=t]}{\mathbb{E}_V[Q_0(1-Q_0)(V,T)|T=t]}\\ &\equiv \frac{\mathbb{E}_V[I(V\le T)\boldsymbol{\psi}(T-V)\mathrm{Var}(J\mid V,T)\mid T=t]}{\mathbb{E}_V[\mathrm{Var}(J\mid V,T)\mid T=t]}\end{aligned}$$

The following proposition, proven in Appendix B.4, establishes that $\hat{\beta}_{K,n}$ is CAN for β_0.

Proposition 1 ($\hat{\beta}_{K,n}$ is CAN)**.** *Assume that Conditions S1-S5 in the Supplement hold. Specifically, let $\alpha_0(t)$ lie in a Hölder class with smoothness index $p = v_0+\gamma$. Let the number of knots $K_n \asymp n^\lambda$. If $1/3 > \lambda$ and $p > 1/2$, then*

$$\sqrt{n}(\hat{\beta}_{K,n}-\beta_0) \rightsquigarrow N(0,\Sigma^{-1})$$

Where $\Sigma \in \mathbb{R}^{d\times d}$ with entries

$$\sigma_{jk} = \mathbb{E}\left[Q_0(V,T)(1-Q_0(V,T))\left(I(V\le T)\psi_j(T-V)-r_{0j}(T)\right)\left(I(V\le T)\psi_k(T-V)-r_{0k}(T)\right)\right]$$

Where Q_0 and r_0 are defined as above.

Inference via Efficient Influence Curve (EIC)

An alternative strategy for inference in 3.5 exploits the theory of influence curves, which can be used to construct asymptotically linear estimators of pathwise differentiable parameters. Popular approaches for semiparametric estimation include one-step estimation[42] and influence-curve-based estimating equations[43,44]. Here, we will discuss targeted maximum likelihood estimation (TMLE)[45], which provides a general framework for using the efficient influence curve (EIC) to construct efficient, robust plug-in estimators with favorable finite-sample properties.

A key step to obtaining the TMLE estimator $\hat{\beta}_{\mathrm{TMLE}}$ of β_0 is the derivation of the EIC. Recall the definitions of $Q_0(V,T)$ and $\boldsymbol{r}_0(T)$ from the previous section. The following Lemma, proven in

Appendix B.5.1, establishes the EIC of β_0 in model 3.5.

Lemma 1 (Efficient Score & EIC of β_0)**.** *The vector-valued efficient Score of β_0 with respect to the semiparametric statistical model in 3.5 is equal to the following:*

$$S_P^{\mathit{eff}}(j,t,v;Q_0,\boldsymbol{r}_0) = \left(I(v \leq t)\boldsymbol{\psi}(t-v) - \boldsymbol{r}_0(t)\right)\left(j - Q_0(v,t)\right)$$

The analogous vector-valued efficient influence function of β_0 is proportional to S^{eff} up to a scaling matrix,

$$D_P(j,t,v;Q_0,\boldsymbol{r}_0) = \Sigma^{-1} S_P^{\mathit{eff}}$$

where the scaling matrix is $\Sigma := \mathbb{E}_0\left[S_P^{\mathit{eff}}(S_P^{\mathit{eff}})^T\right]$.

At a high-level, the TMLE algorithm works in three steps (see Appenix B.5.2 for details). In step one, we obtain initial estimates of the relevant components of the distribution P_0 (Q_0 and r_0 in this case) which we denote $P_n^0 \equiv (\hat{Q}_n^0, \hat{r}_n^0)$. In step two, calculate a "clever covariate", denoted $H := (I(v \leq t)\boldsymbol{\psi}(t-v) - \boldsymbol{r}_0(t))$, by plugging in our pilot estimator of $\boldsymbol{r}_0$. Denote the initial estimator of the clever covariate $\hat{H}^0$. In step three, starting at $k = 0$, we iteratively update the nuisance parameter fits P_n^{k+1} using a working logistic model for J with main effect $\hat{H}^k$ and offset $\text{logit}(\hat{Q}_n^k)$ until the nuisance updates (i.e., coefficients on the clever covariate) are small. The inclusion of the clever covariate ensures the logistic model "targets" the relevant components of the data distribution $P_n^* = (\hat{Q}_n^*, \hat{r}_n^*)$ to solve the estimating equation associated with the EIC. According to the theory of TMLE, under regularity conditions, the plug-in estimator using the targeted components P_n^* achieves the optimal bias-variance tradeoff and will be CAN for the target parameter, β_0, with asymptotic variance characterized by the variance of the EIC. The following proposition, proven in Appendix B.5.3 establishes that the TMLE estimator is CAN under some regularity conditions.

Proposition 2 ($\hat{\beta}_{\text{TMLE}}$ is CAN)**.** *Suppose Conditions S6-S9 hold. In particular, we require that our estimators of the relevant components of the EIC are consistent with the following rates $||\hat{r}_n^* - r_0|| =$*

$o_P(n^{-1/4})$ *and* $||\hat{Q}_n^* - Q_0|| = o_P(n^{-1/4})$ *where* $||\cdot||$ *refers to the* L_P^2 *norm. Then*

$$\sqrt{n}(\hat{\beta}_{\text{TMLE}} - \beta_0) \rightsquigarrow N(0, \text{Cov}(D_P))$$

Where the convergence in distribution is to a random vector in $\mathbb{R}^d$ *and*

$$\text{Cov}(D_P) := \mathbb{E}[D_P(J,T,V;Q_0,\boldsymbol{r}_0)(D_P(J,T,V;Q_0,\boldsymbol{r}_0))^T] \equiv \Sigma^{-1}$$

is the $d \times d$ *covariance matrix of the EIC.*

Note that under the appropriate regularity conditions on each estimator, the asymptotic variance of the TMLE estimator is equivalent to the asymptotic variance of the sieve estimator. This underscores the efficiency of both estimators.

3.2.6 Extensions

Enforcing monotonicity of time-varying VE

For many vaccines, outside a short period of time for the vaccine to produce an immune response, it is plausible that VE may remain constant or decrease, but cannot increase in time-since-vaccination. To produce sensible estimators which agree with scientific knowledge and improve finite sample performance, we may wish to enforce $\text{VE}_1(\tau)$ to be monotonically non-increasing. As explained in Appendix B.8.1, following results by Westling et al. [46], we can obtain a corrected monotonized estimator (and associated confidence interval limits), by running a precision-weighted isotonic regression smoothing of the unconstrained estimator $\hat{f}$ and its confidence limits over a fine grid of values of time-since-vaccination. Westling et al. [46] established that the estimation error of the monotonized estimator over the grid is never worse than the unconstrained estimator, that the monotonized confidence limits do not degrade coverage or bandwidth over the grid relative to the unconstrained estimator's limits.

Extension to strain-specific, time-varying VE

Since COVE conducted whole genome sequencing of swabs positive for SARS-CoV-2, we can tweak our approach to estimate *strain-specific, time-varying VE.*[47] We refer the reader to Appendix

B.8.3 for full details. We recast our semiparametric *binary* logistic regression model for vaccine-preventable infection ($J = 1$) as a semiparametric *multinomial* logistic regression model for vaccine-preventable infection ($J = 1$) with strain $S \in \{1, \ldots, M\}$. We can write the multinomial logistic model equivalently as a series of M independent binary regressions of each vaccine-preventable infection strain against the irrelevant infection as the pivot outcome:

$$\log \left\{ \frac{P(J = 1, S = s \mid V, T = t)}{P(J = 0 \mid V, T = t)} \right\} = I(V \leq t)\beta_{0s}^T \boldsymbol{\psi}_s(t - v) + \alpha_0(t) + \log(p_{0s}(t)). \tag{3.7}$$

Time-varying VE against the vaccine-preventable pathogen of strain s is given by $\text{VE}_s(t - v) = 1 - \exp\{\beta_{0s}^T \boldsymbol{\psi}_s(t - v)\}$, where β_{0s} is a d-dimensional unknown Euclidean parameter and $\boldsymbol{\psi}_s$ is a known basis expansion of time-since-vaccination. As before, $\alpha_0(t) := \log\{P(J = 1 \mid V > t, T = t)/P(J = 0 \mid V > t, T = t)\}$ reflects the relative background calendar time trends of vaccine-preventable and irrelevant infections. The model also depends on M variant mixture proportions $\{p_{0s'}(t)\}_{s=1}^M$ such that $p_{0s'}(t) := P(S = s' \mid T = t, J = 1, V > t)$ which model the background probability of acquiring vaccine-preventable infection with strain $S = s'$ conditional on a vaccine-preventable infection of unknown strain at time t. In our case study, we used external, population-level data to profile out the mixture proportions and improve model identifiability, a technique referred to as surveillance anchored sieve analysis.[30] We estimated the mixture proportions using weekly summaries of SARS-CoV-2 variant proportions in the US from GISAID.[48] In the time period spanning COVE's booster phase, GISAID estimates were based on a weekly average of roughly 62,500 genome submissions and Delta and Omicron were predominant with minimal overlap, hence, we treat the mixture proportions as fixed and known.

With known variant mixture proportions, we can directly extend the sieve estimator from the binary case to the multinomial case. Specifically, we maximize the multinomial sieve likelihood with respect to the stacked parameter $\beta_0 := (\beta_{0s=1}, \ldots, \beta_{0s=M})$ and the coefficients associated with $\alpha_0(t)$, which is approximated with a sieve (see Appendix B.8.3 for details). Consistency and asymptotic normality follow from the same arguments used in the binary case. Extending the TMLE approach requires some minor modifications, including deriving a new EIC and clever covariate. Associated derivations and a description of the TMLE algorithm adapted to the multinomial case can be found in Appendix B.8.3.

3.3 Connection to a causal VE waning parameter

In this section, we further study the causal interpretation of our hazard-based statistical parameter under appropriate assumptions, as others have done with hazard-based contrasts.[10,49,50]

We consider potential infection outcomes that depend on a joint intervention that defers vaccination and exposure until specific dates. While administration of a particular vaccine regimen (e.g., single-dose mRNA-1273) on a specific date is a clear and well-defined intervention, we must carefully define "exposure" and consider what it means to intervene upon exposure. One possible definition of "exposure" is *virological*: a controlled administration (e.g., via intranasal pipette) of a known quantity of virus particles capable of causing infection in a naïve patient.[13] One can conceptualize intervening on virological exposure in a challenge trial where the timing of challenge exposure is under the investigator's control. However, applying the virological definition of exposure to real-world studies is complicated by differences between idealized "challenge" exposures and real-world exposures due to changing viral strains and variable viral loads in the population. Instead, we adopt an *epidemiological* definition of exposure[51] where a participant forgoes a period of strict isolation and may interact with the community in which pathogens circulate. While it is difficult and perhaps unethical to isolate participants from the community, one can conceptualize a trial that randomizes participants to receive vaccination and cease isolating themselves on different dates. Trials that randomize participants to receive a vaccine on different dates before the flu season approximate deferred exposure designs, as participants are essentially "isolated" when influenza circulation is low.[9,12]

Let the time origin $t_0 = 0$ refer to the study start date. Let $T_1(v, d)$ denote the hypothetical time-to-vaccine-preventable infection under an intervention which defers vaccination until date $V = v \geq 0$ and epidemiological exposure (i.e., interaction with the community) until date $d \geq 0$. Let $T_0(v, d)$ denote the analogous potential vaccine-irrelevant infection time. Let $T(v, d)$ denote $\min\{T_1(v, d), T_0(v, d)\}$. We define the following causal time-varying VE parameter.

$$\mathrm{VE}_C(t - v) = 1 - \theta_C(t - v) = \lim_{\epsilon \downarrow 0} 1 - \frac{P(T_1(v, 0) \in [t, t + \epsilon) \mid T(v, 0) \geq t, V = v)}{P(T_1(\infty, t) \in [t, t + \epsilon) \mid T(v, 0) \geq t, V = v)} \tag{3.8}$$

Consider a target trial in which all individuals are vaccinated at time $V = v$ and are permitted to interact with the community from study entry onward. Among those who remain uninfected at

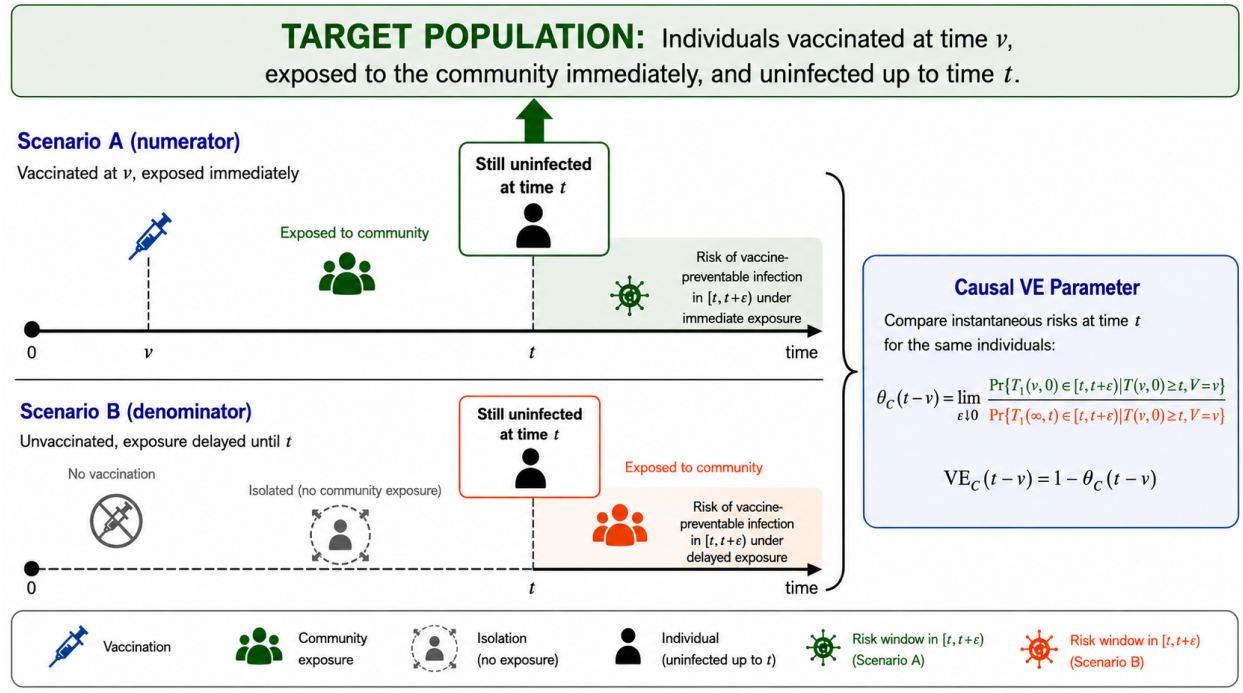

Figure 3.1: Figure illustrating interpretation of causal time-varying VE parameter.

time t, the parameter (3.8) compares counterfactual rates of vaccine-preventable infection under (i) immediate exposure following vaccination, and (ii) the risk that would have been observed had these same individuals not been vaccinated but instead avoided exposure by remaining isolated from the community until time t. Unlike typical hazard-based contrasts, both quantities are evaluated among the same population. Hence, (3.8) admits a direct causal interpretation of a well-defined intervention in a single, observed subgroup of participants.

Intuitively, the parameter addresses the question: *"among individuals vaccinated at time v and exposed immediately who remain uninfected at time t, how does their risk of infection at time t compare to the risk they would have experienced had they remained unvaccinated but avoided exposure by isolating until time t?"* We illustrate this interpretation graphically in Figure 3.1.

To identify 3.8, we assume the hazard ratio of vaccine-preventable and irrelevant infections is stable across subgroups when the unvaccinated, deferred exposure regime is enforced.

Assumption 9 (Bias equivalence)**.** *Suppose that for all* $t > 0$ *and* $v > 0$

$$\lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty, t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)}{P(T_1(\infty, t) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)} = \lim_{\epsilon \downarrow 0} \frac{P(T_0(\infty, t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)}{P(T_0(\infty, t) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)}.$$

Equivalently, the ratio of cause-specific hazards under a hypothetical regime deferring vaccination and exposure is stable across subgroups of participants defined by principal stratum membership. That is, for all $t > 0$*, and* $v > 0$*,*

$$\lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty, t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)}{P(T_0(\infty, t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)} = \lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty, t) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)}{P(T_0(\infty, t) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)} = \alpha_0(t)$$

Assumption 9 claims that the bias (i.e., quantified by the shift in distribution of vaccine-free potential outcomes between subgroups) inflicted on the hazard ratio of vaccine-preventable infection is equivalent to the bias inflicted on hazard ratio of vaccine-irrelevant infections. We can justify the Assumption 9 using an analogy. Suppose that two groups play dodgeball for ϵ minutes under identical conditions. Players are "out" when they are hit by a red or blue dodgeball, and all players are eligible to be "out" at the start of the playing period. Suppose red balls are twice as abundant as blue balls in the gym. While one group may have better dodging skills than the other group, it is reasonable the number of players hit by a red ball should be twice the number of players hit by a blue ball in both groups. We adapt this analogy to our setting by letting red and blue balls represent vaccine-preventable and vaccine-irrelevant infections, recasting the time-varying ratio of red to blue balls in the gym to $\alpha_0(t)$, changing the two groups of players to the two principal subgroups ($\{T(v,0) \geq t, V = v\}$ and $\{T(\infty,0) \geq t, V = \infty\}$), recasting the game duration as the exposure increment ϵ. Importantly, Assumption 9 does not require that the infection rates be identical for different subgroups. Instead, Assumption 9 requires the *ratio* of vaccine-preventable to vaccine-irrelevant infections is stable across the different groups.

Our next assumption requires that the hazard of vaccine-irrelevant infections does not depend on the vaccination date.

Assumption 10 (NCO)**.** *Assume that the vaccine-irrelevant infection is unaffected by vaccination date, such that for all* $t > 0$*,* $v > 0$*, and any* $a, a' > 0$ *such that* $a \neq a'$*,*

$$P(T_0(a, t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v) = P(T_0(a', t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)$$

We believe this assumption is plausible in settings where the vaccine is unlikely to affect the

NCO, which is most likely to hold when the vaccine is highly antigenically specific and the NCO is sufficiently distinct from the vaccine's target. The lack of an observed vaccine effect on non-SARS ARIs in COVE's blinded phase,[29] provides empirical support for this assumption in our case study.

In real-world studies, all participants interact with the community at study entry, and we never observe data under the "deferred exposure" world. However, it may be reasonable that among participants who were *naturally* uninfected by time t, their failure rates just after time t would be unchanged had they been isolated until t. Assumption 11 provides the crucial link between the counterfactual "deferred exposure" world and the factual "immediate exposure" world.

Assumption 11 (Exposure deferral irrelevance)**.** *for both* $j \in \{0,1\}$ *and any* $t > 0$*,*

$$P(T_j(v,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v) = P(T_j(v,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)$$

Assumption 11 is a strong assumption analogous to Assumption 19 in Janvin and Stensrud[13] that contends that being naturally uninfected by time t simulates isolation from the community until t. Assumption 11 could be violated if there are a significant number of sub-infectious pathogen exposures or asymptomatic infections during the period from 0 to t that reduce infection risk at t. Assumption 11 may be plausible if sub-infectious pathogen exposures or asymptomatic infections are rare or are unlikely to influence subsequent infection risk, as participants can be effectively considered isolated.[9,12]

As shown in Appendix B.3, under the above assumptions and consistency of potential outcomes, the causal quantity $\theta_C(t-v)$ is identified by the ratio of cause-specific hazard ratios. As before, applying conditional probability rules shows that the causal parameter is also identified by the odds ratio of infection causes at time t.

$$\theta_C(t-v) = \frac{P(J=1 \mid T=t, V=v)/P(J=1 \mid T=t, V=\infty)}{P(J=0 \mid T=t, V=v)/P(J=0 \mid T=t, V=\infty)}$$

In Section 1.3.2 of the supplementary materials, we introduce a sensitivity analysis to quantify robustness of our results to deviations from Assumptions 9-11 by analyst-defined constants.

3.4 Case Study: Time-varying effectiveness of mRNA-1273 booster in an observational extension of a COVID-19 vaccine trial

The COVE booster analysis set contained 20,404 participants who previously received two doses of mRNA-1273 vaccine (either at randomization or open-label crossover) without virological/serologic evidence of SARS-CoV-2 infection on or prior to September 23, 2021.[27] Therefore all mRNA-1273 booster VE estimates are relative VE estimates for COVID-19 disease comparing 3 doses (primary series + booster) to 2 doses (primary series only). In response to acute respiratory illness (ARI) symptoms, two nasopharyngeal swabs were collected. One was tested for SARS-CoV-2 using RT-PCR and the other was tested for 20 other respiratory pathogens using a multiplex PCR assay (Biofire® RP2.1). COVE also conducted whole genome sequencing of SARS-CoV-2 positive swabs to identify the infecting strain. In COVE's booster phase, 441 of 1902 (23%) of swabs positive for SARS-CoV-2 had a missing lineage group. In the US in Winter 2021, Delta was the exclusive circulating lineage for months before being rapidly and fully displaced by Omicron, with only a short period of co-circulation. Hence, lineage could be viewed as a near deterministic function of calendar time, and imputation uncertainty was almost negligible. Hence, for the strain-specific analysis, missing SARS-CoV-2 lineage groups (Delta or Omicron) were singly imputed by sampling from a binomial distribution using weekly variant mixing proportions obtained from GISAID.[48]

3.4.1 Time-varying VE during periods where Delta and Omicron were predominant

We estimated time-varying VE of the mRNA-1273 booster in COVE for (a) the Delta period (September 23, 2021 – December 11, 2021) and (b) the Omicron period (December 18, 2021 – April 05, 2022). We fit the proposed SLR estimators using Sieve and TMLE techniques and a Cox regression estimator for comparison. All estimators used a single parameter to model VE [0,13] days post-boost and allowed VE to vary linearly on the log hazard scale starting 14 days post-boost. For the sieve estimator, we approximated $\alpha_0(t)$ using a B-spline basis of dimension $K = \lfloor n^{1/3.5} \rfloor$ to satisfy the knot growth constraint in Proposition 1. For the TMLE estimator, $\alpha_0(t)$ and $\mathbf{r}_0(t)$ were estimated using generalized additive models (GAMs) using the default behavior in 'mgcv::s()': a thin plate regression spline basis of dimension 10 with automatic smoothing parameter selection. To produce sensible estimators, we applied the precision-weighted isotonic correction to

the unconstrained estimates of f as described in Subsection 3.2.6. For comparison, fit a Cox model ignoring vaccine-irrelevant infections and parametrized as follows

$$\lambda_S(t \mid V, X) = \lambda_0^S(t) \exp\{Z(t)I(\tau \leq 13)\beta_1 + Z(t)I(\tau > 13)\beta_2(\tau - 13) + X^\top \gamma\}$$

Where S referred to the membership in a stratum defined by risk of severe COVID-19, sex, and underrepresented racial/ethnic minority status, $Z(t) := I(V \leq t)$ denoted current vaccination status, $\tau := \max(0, V - t)$ denoted time-since-vaccination, and X referred to a continuous baseline risk score previously developed to predict COVID-19 risk in the placebo arm using demographic and clinical characteristics (see Gilbert et al. [52] for details).

Analysis results are shown in Figure 3.2. During the Delta period, Cox regression reported high initial efficacy (VE at 14 days post-boost = 87.5%; 95% CI: 67.2% to 95.2%) and durable protection 2 months after the boost (VE at 60 days = 81.0%; 95% CI: 50.5% to 92.7%). Compared to Cox regression, the TMLE estimator supported slightly higher initial effectiveness (VE at 14 days post-boost = 91.8%; 95% CI: 83.7% to 95.9%) and more pronounced waning over 2 months (VE at 60 days = 72.6%; 95% CI: -11.1% to 93.2%). Accounting for vaccine-irrelevant ARIs led to substantially different results in the Omicron period. Cox regression suggested moderate initial effectiveness (VE at 14-days post-boost = 58.9%; 95% CI: 48.6% to 67.1%) which rapidly decayed over 120 days (VE at 4 months post-boost = 14.0%; 95% CI: -2.5% to 27.9%). In contrast, the TMLE estimator supported higher initial effectiveness (VE at 14-days post-boost = 80.1%; 95% CI: 57.4% to 90.7%) and more durable VE (VE at 4 months post-boost = 60.9%; 95% CI: 22.8% to 80.2%). Results for the Sieve estimator were comparable to the TMLE estimator for both Delta and Omicron period analyses.

3.4.2 Strain-specific time-varying VE against Delta and Omicron COVID-19

Since COVE conducted whole genome sequencing of swabs from COVID-19 cases, we can estimate *strain-specific, time-varying VE* by slightly modifying our approach as described in 3.2.6. Strain-specific VE estimates improve upon the period-specific estimates by avoiding strain contamination and extending VE estimation to periods of Delta and Omicron co-circulation. For our proposed estimators, we assumed the variant mixture proportions were known and equaled those provided

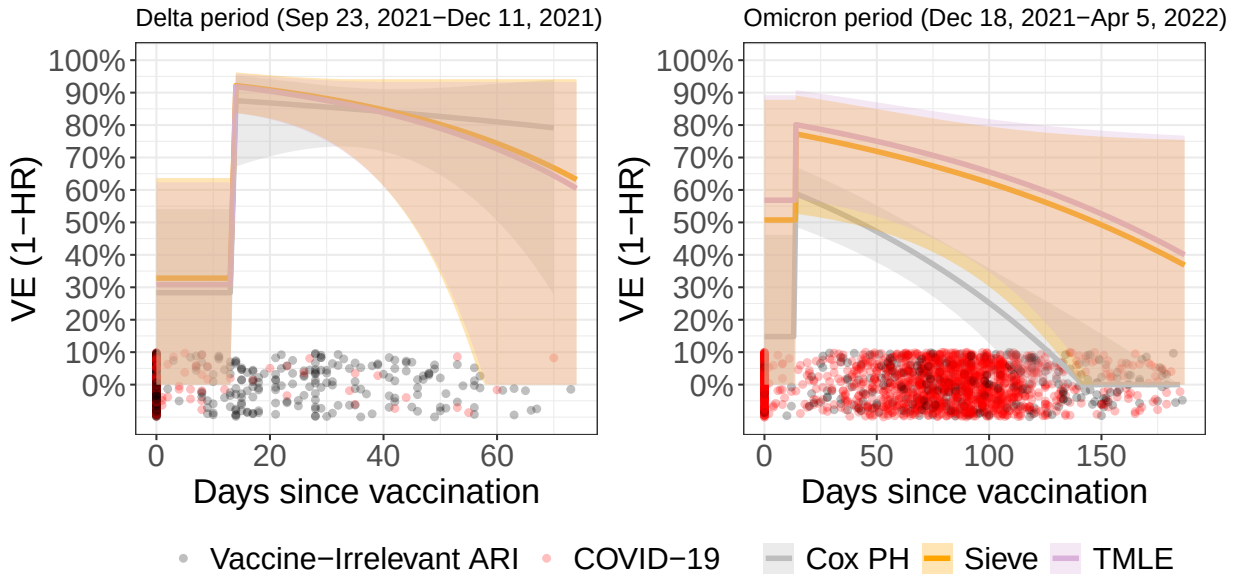

Figure 3.2: Summary of single mRNA-1273 booster efficacy in Delta (left) and Omicron period (right). Time-varying VE estimated using semiparametric logistic regression model with Sieve and TMLE estimators and Cox regression. Vaccine-irrelevant and COVID-19 illnesses shown as points on bottom of plot. Infections pre-vaccination stacked at $\tau = 0$.

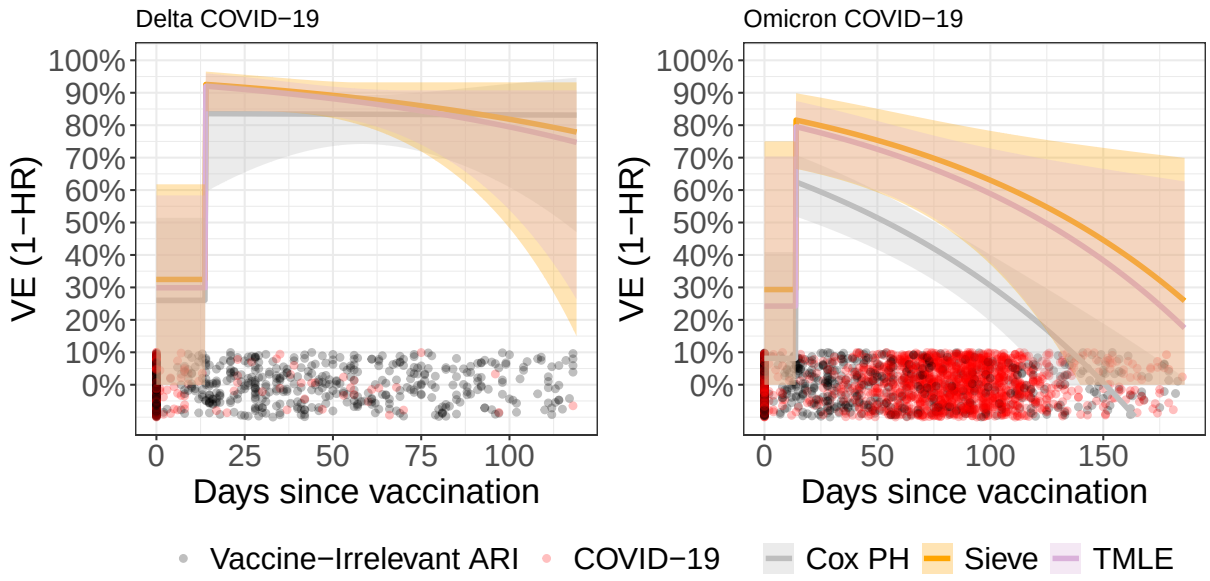

Figure 3.3: Summary of single mRNA-1273 booster efficacy in Delta (left) and Omicron period (right). Time-varying VE estimated using semiparametric logistic regression model with Sieve and TMLE estimators, and standard Cox PH regression. Vaccine-irrelevant and COVID-19 infections shown as points on bottom of plot. Infections pre-vaccination stacked at $\tau = 0$.

by GISAID.[48] We fit a semiparametric multinomial logistic regression model using both sieve and TMLE approaches. For the sieve estimator, we approximated $\alpha_0(t)$ using a B-spline basis with 13 interior knots. The TMLE estimator used kernel estimators for the unknown nuisance parameters with an epanechnikov shape and bandwidth chosen using Silverman's rule of thumb. All our proposed estimators assumed a single scalar modeling VE [0,13] days post-boost and assumed a linear specification for $f(\cdot)$ thereafter. For comparison, we fit Cox regression models to Delta and Omicron infections, where the occurrence of the other outcome was treated as censored by a competing risk.[47] As before, the Cox regression analyses ignored data on vaccine-irrelevant ARIs but stratified by randomization stratum, sex, and underrepresented racial/ethnic minority status and adjusted for a continuous baseline risk score developed using ensemble machine learning[52].

Results can be found in Figure 3.3. Cox regression which ignored vaccine-irrelevant ARIs supported high initial booster VE against Delta COVID-19 (VE at 14 days post-boost = 83.5%; 95% CI: 59.4% to 93.3%) that remained constant for 100 days (VE at 100 days post-boost = 83.2%;

95% CI: 60.9%, 92.8%). Compared to Cox regression, the TMLE estimator accounting for vaccine-irrelevant ARIs suggested higher initial VE against Delta (VE at 14 days post-boost = 92.0%; 95% CI: 84.3% to 96.0%) and more marked VE waning (VE at 100 days post-boost = 79.3%; 95% CI: 53.5%, 90.7%). Against Omicron COVID-19, Cox regression suggested moderate booster effectiveness (VE at 14 days post-boost = 62.5%; 95% CI: 51.8%, 70.8%) that waned substantially over 5 months (VE at 150 days post-boost = 0.6%; 95% CI: -21.7%, 18.9%). In contrast, our proposed TMLE estimator supported that the mRNA-1273 booster was more effective (VE at 14 days post-boost = 79.3%; 95% CI: 66.3%, 87.3%) and more durable (VE at 150 days post-boost = 38.3%; 95% CI: -13.8%, 66.5%) against Omicron COVID-19. In both Delta and Omicron analyses, the sieve estimator results were comparable to the TMLE estimator.

3.5 Discussion

Herein, we present a framework to eliminate unmeasured confounding and selection bias from hazard-based, time-varying vaccine effectiveness (VE) estimates by using vaccine-irrelevant infections as negative controls. We used this method to explore the effectiveness of a single mRNA-1273 booster in an observational extension of a COVID-19 vaccine trial that spanned the Omicron COVID-19 wave and used multiplex PCR testing to detect acute respiratory illnesses (ARIs) caused by SARS-CoV-2 and 20 other "off-target" respiratory pathogens. Accounting for the vaccine-irrelevant ARIs supported waning vaccine effectiveness against Delta COVID-19 and more effective and durable protection against Omicron COVID-19 than suggested by Cox regression analyses ignoring vaccine-irrelevant ARIs.

At first glance, it may seem puzzling why accounting for vaccine-irrelevant ARIs seemingly had a greater impact on time-varying VE estimates for the Omicron period/variant analyses compared to the Delta period/variant analyses. One hypothesis is that the higher force of infection during the Omicron wave accelerated depletion of susceptibles,[11,14] a phenomenon documented in our simulations. If depletion of susceptibles was higher during in the Omicron period/variant analyses, and vaccine-irrelevant infections were a good proxy for the imbalance in infection risk, we would anticipate a larger bias correction in the Omicron analyses.

We acknowledge some important caveats when interpreting our case study results. First, COVID-19 vaccine effectiveness may depend on viral sequence marks within a single lineage,[53,54]

which may lead to minor violations of Assumption 3 if within-lineage viral sequence characteristics change over calendar time. Adapting our approach to accommodate continuous viral marks is an area of future work. Second, our developments assumed a participant's infection outcomes do not depend on the vaccination statuses of other study participants or community members. Since study participants were embedded in a much larger population, interactions between study participants were very unlikely and lack of interference between study participants is likely. However, a study participant's infection outcomes may depend on the vaccination statuses of their community contacts. We believe our statistical parameter $\text{VE}_1(t-v)$ could be robust to community interference; following the logic of Perényi and Stensrud[55], since our statistical parameter conditions on infection and by extension, exposure to the infected individual, our parameter could be robust to the particular contact/interference structure in the population.[55] Exploring our estimator's robustness to interference is an interesting avenue for further research.

We anticipate that the approach described here could be efficiently applied to case-only study designs, such as the popular test negative design (TND) – a variant of a case-control design which conditions on seeking testing in response to acute respiratory illness. In a TND, it is common to define symptomatic illness with a negative test test for the pathogen of interest as the "control" outcome, which is assumed to be unaffected by vaccination. A recent meta-analysis identified that test-negative ARIs may be compelling choices of NCOs, although care needs to be taken in rigorously defining a test negative event.[56] Acute respiratory illness with a *positive* test for an unrelated pathogen may be a more objective NCO endpoint by avoiding arbitrary aspects of test-negative event definitions and may exhibit better overlap in unmeasured causes with the primary infection endpoint.[29] Establishing assumptions under which our approach yields reliable estimates of time-varying VE in TNDs is a promising direction for future work. Any future efforts to infer time-varying VE from TNDs should endeavor to account for correlated uptake of vaccines against several vaccine-preventable pathogens (e.g., COVID-19, influenza, RSV, etc.), natural immunity caused by prior infection, and imperfect outcome ascertainment. Extending recent double-negative control approaches to estimate time-varying VE in TNDs is another avenue for future work.[57]

Moreover, despite being a popular method among practitioners, the causal interpretation of hazard-based treatment effect estimates remains unclear in general.[10,49,58] We propose a time-varying, hazard-based statistical parameter that adjusts for unmeasured sources of bias, and en-

deavor to describe its causal interpretation under some assumptions. Recent work by Janvin and Stensrud [13] defined a causal waning VE parameter motivated by challenge trials and proposed assumptions to bounds to bound their unknown parameter. The potential of using irrelevant infections to sharpen partial identification results and motivate alternative causal VE waning parameters is a compelling area of future work.

Ethics and Data Availability Statement

The study uses de-identified data from the COVE study. The study sponsor, Moderna, Inc., was responsible for conceptualization, trial design, site selection, & monitoring, The trial was conducted in accordance with the International Council for Harmonisation of Technical Requirements for Registration of Pharmaceuticals for Human Use, Good Clinical Practice guidelines. A central Institutional Review Board (Advarra, Inc., Columbia, MD) approved the protocol and consent forms. All participants provided written informed consent.

Code developed for simulations and case study are publicly available on GitHub. Access to participant-level data and supporting clinical documents by qualified external researchers may be made available upon request and subject to review. A materials transfer and/or data access agreement with the sponsor will be required for accessing shared data.

3.6 Acknowledgments and Funding Statement

We gratefully acknowledge Craig Magaret and all data contributors, authors, and laboratories for generating the genetic sequence and metadata and sharing via the GISAID Initiative. The study sponsor, Moderna, Inc., was responsible for conceptualization, trial design, site selection, and monitoring. Statistical analysis was supported by the National Science Foundation Graduate Research Fellowship Program (Grant No. DGE-2140004). Any opinions, findings, conclusions, or recommendations expressed in this material do not necessarily reflect the views of the National Science Foundation. We also would like to acknowledge the NIH R01AI192632 awarded to BZ, HJ, PBG, HES.

Bibliography

[1] Edward A. Belongia, Maria E. Sundaram, David L. McClure, Jennifer K. Meece, et al. Waning vaccine protection against influenza a (h3n2) illness in children and older adults during a single season. *Vaccine*, 33(1):246–251, 2015. ISSN 0264-410X. doi: https://doi.org/10.1016/j.vaccine.2014.06.052. URL `https://www.sciencedirect.com/science/article/pii/S0264410X14008536`.

[2] Jill M. Ferdinands, Alicia M. Fry, Sue Reynolds, Joshua G. Petrie, et al. Intraseason waning of influenza vaccine protection: Evidence from the us influenza vaccine effectiveness network, 2011–2012 through 2014–2015. *Clinical Infectious Diseases*, 64(5):544–550, 02 2017. ISSN 1058-4838. doi: 10.1093/cid/ciw816. URL `https://doi.org/10.1093/cid/ciw816`.

[3] Dan-Yu Lin, Donglin Zeng, and Peter B Gilbert. Evaluating the long-term efficacy of coronavirus disease 2019 (COVID-19) vaccines. *Clin. Infect. Dis.*, 73(10):1927–1939, November 2021.

[4] Jonathan Fintzi and Dean Follmann. Assessing vaccine durability in randomized trials following placebo crossover. *Statistics in Medicine*, 40(27):5983–6007, 2021.

[5] Dan-Yu Lin, Yu Gu, Bradford Wheeler, Hayley Young, et al. Effectiveness of Covid-19 Vaccines over a 9-Month Period in North Carolina. *New England Journal of Medicine*, 386(10):933–941, March 2022. ISSN 0028-4793, 1533-4406. doi: 10.1056/NEJMoa2117128. URL `http://www.nejm.org/doi/10.1056/NEJMoa2117128`.

[6] Dan-Yu Lin, Yangjianchen Xu, Yu Gu, Donglin Zeng, et al. Durability of Bivalent Boosters against Omicron Subvariants. *New England Journal of Medicine*, 388(19):1818–1820, May 2023. ISSN 0028-4793, 1533-4406. doi: 10.1056/NEJMc2302462. URL `http://www.nejm.org/doi/10.1056/NEJMc2302462`.

[7] Dan-Yu Lin, Yi Du, Yangjianchen Xu, Sai Paritala, Matthew Donahue, and Patrick Maloney. Durability of XBB.1.5 vaccines against omicron subvariants. *New England Journal of Medicine*, 390(22):2124–2127, June 2024.

[8] Mark Lipsitch. Challenges of vaccine effectiveness and waning studies. *Clinical Infectious Diseases*, 68, 2018.

[9] M. Lipsitch, E. Goldstein, G.T. Ray, and B. Fireman. Depletion-of-susceptibles bias in influenza vaccine waning studies: how to ensure robust results. *Epidemiology & Infection*, 2019.

[10] Miguel A Hernán. The hazards of hazard ratios. *Epidemiology*, 21(1):13–15, 2010.

[11] Rebecca Kahn, Stephanie J Schrag, Jennifer R Verani, and Marc Lipsitch. Identifying and Alleviating Bias Due to Differential Depletion of Susceptible People in Postmarketing Evaluations of COVID-19 Vaccines. *American Journal of Epidemiology*, 191(5):800–811, March 2022. ISSN 0002-9262, 1476-6256. doi: 10.1093/aje/kwac015. URL `https://academic.oup.com/aje/article/191/5/800/6515669`.

[12] Thomas Ray, Ned Lewis, Nicola P Klein, Matthew F Daley, et al. Depletion-of-susceptibles Bias in Analyses of Intra-season Waning of Influenza Vaccine Effectiveness. *Clinical Infectious Diseases*, 70(7):1484–1486, March 2020. ISSN 1058-4838, 1537-6591. doi: 10.1093/cid/ciz706. URL `https://academic.oup.com/cid/article/70/7/1484/5540005`.

[13] Matias Janvin and Mats J Stensrud. Quantification of vaccine waning as a challenge effect. *J. Am. Stat. Assoc.*, pages 1–11, December 2024.

[14] Mona N. Kanaan and C. Paddy Farrington. Estimation of Waning Vaccine Efficacy. *Journal of the American Statistical Association*, 97(458):389–397, 2002. ISSN 0162-1459. URL `https://www.jstor.org/stable/3085652`. Publisher: [American Statistical Association, Taylor & Francis, Ltd.].

[15] World Health Organization. Evaluation of COVID-19 vaccine effectiveness: Interim Guidance, March 2021.

[16] World Health Organization. Evaluation of COVID-19 vaccine effectiveness in a changing landscape of COVID-19 epidemiology and vaccination: Interim Guidance, October 2022.

[17] M.E. Halloran, I.M. Longini, and C.J. Struchiner. Estimability and interpretation of vaccine efficacy using frailty mixing models. *American Journal of Epidemiology*, 144:83–97, 1996.

[18] Anastasios Tsiatis and Marie Davidian. Estimating vaccine efficacy over time after a randomized study is unblinded. *Biometrics*, 78(3):825–838, 2021.

[19] Lisa A Jackson, Michael L Jackson, Jennifer C Nelson, Kathleen M Neuzil, and Noel S Weiss. Evidence of bias in estimates of influenza vaccine effectiveness in seniors. *International Journal of Epidemiology*, 35:337–344, 2005.

[20] Michael L Jackson and Jennifer C Nelson. The test-negative design for estimating influenza vaccine effectiveness. *Vaccine*, 31(17):2165–2168, 2013.

[21] Korryn Bodner, Michael A. Irvine, Jeffrey C. Kwong, and Sharmistha Mishra. Observed negative vaccine effectiveness could be the canary in the coal mine for biases in observational COVID-19 studies. *International Journal of Infectious Diseases*, 131:111–114, June 2023. ISSN 12019712. doi: 10.1016/j.ijid.2023.03.022. URL `https://linkinghub.elsevier.com/retrieve/pii/S1201971223001029`.

[22] C.P. Farrington, M.N. Kanaan, and N.J. Gay. Estimation of the basic reproduction number for infectious diseases from age-stratified serological survey data. *Applied Statistics*, 50:251–292, 2001.

[23] C Paddy Farrington, Heather J Whitaker, Steffen Unkel, and Richard Pebody. Correlated infections: quantifying individual heterogeneity in the spread of infectious diseases. *Am. J. Epidemiol.*, 177(5):474–486, March 2013.

[24] Steffen Unkel, C. Paddy Farrington, Heather J, Whitaker, and Richard Peabody. Time varying frailty models and the estimation of heterogeneities in transmission of infectious diseases. *Journal of the Royal Statistical Society: Series C*, 63(1):141–158, 2014.

[25] Lindsey R Baden, Hana M El Sahly, Brandon Essink, Karen Kotloff, et al. Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine. *New England Journal of Medicine*, 384(5):403–416, 2021.

[26] Hana M El Sahly, Lindsey R Baden, Brandon Essink, Susanne Doblecki-Lewis, et al. Efficacy of the mrna-1273 sars-cov-2 vaccine at completion of blinded phase. *New England Journal of Medicine*, 385(19):1774–1785, 2021.

[27] Lindsey R. Baden, Hana M. El Sahly, Brandon Essink, Dean Follmann, et al. Long-term safety and effectiveness of mRNA-1273 vaccine in adults: COVE trial open-label and booster phases. *Nature Communications*, 15(1):7469, August 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-50376-z. URL `https://www.nature.com/articles/s41467-024-50376-z`.

[28] M Elizabeth Halloran, Michael Haber, Ira M Longini, and Claudio J Struchiner. Direct and indirect effects in vaccine efficacy and effectiveness. *American Journal of Epidemiology*, 133 (4):323–331, 1991.

[29] Ethan Ashby, Holly Janes, Dean Follmann, Peter B. Gilbert, et al. Validating and leveraging non-sars-cov-2 respiratory infection as a negative control outcome in a phase 3 covid-19 vaccine trial with extended observational follow-up. *The American Journal of Epidemiology*, 2025. https://doi.org/10.1093/aje/kwaf176.

[30] Dean Follmann, Michael Fay, and Craig Magaret. Estimation of vaccine efficacy for variants that emerge after the placebo group is vaccinated. *Statistics in Medicine*, 41(16):3076–3089, July 2022.

[31] Marc Lipsitch, Eric Tchetgen Tchetgen, and Ted Cohen. Negative Controls: A Tool for Detecting Confounding and Bias in Observational Studies. *Epidemiology (Cambridge, Mass.)*, 21(3):383–388, May 2010. ISSN 1044-3983. doi: 10.1097/EDE.0b013e3181d61eeb. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3053408/`.

[32] Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. A Selective Review of Negative Control Methods in Epidemiology. *Current Epidemiology Reports*, 7(4):190–202, December 2020. ISSN 2196-2995. doi: 10.1007/s40471-020-00243-4. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8118596/`.

[33] Eric J. Tchetgen Tchetgen, Tamar Sofer, and David Richardson. Negative outcome control for unobserved confounding under a cox proportional hazards model. Technical report, Harvard University, 2015.

[34] Eric J. Tchetgen Tchetgen, Chan Park, and David B. Richardson. Universal difference-in-differences for causal inference in epidemiology. *Epidemiology*, 35, 2024.

[35] Christopher B Boyer, Kendrick Qijun Li, Xu Shi, and Eric J Tchetgen Tchetgen. Identification and estimation of vaccine effectiveness in the test-negative design under equi-confounding. *Epidemiology*, 37(1):77–87, January 2026.

[36] Peter B Gilbert, L J Wei, MR Kosorok, and JD Clemens. Simultaneous inferences on the contrast of two hazard functions with censored observations. *Biometrics*, 58:773–80, 2002.

[37] Eric J Tchetgen Tchetgen, James M Robins, and Andrea Rotnitzky. On doubly robust estimation in a semiparametric odds ratio model. *Biometrika*, 97(1):171–180, March 2010.

[38] Zhiqiang Tan. On doubly robust estimation for logistic partially linear models. *Statistics & Probability Letters*, 155:108577, December 2019. ISSN 0167-7152. doi: 10.1016/j.spl.2019.108577. URL `https://www.sciencedirect.com/science/article/pii/S0167715219302238`.

[39] Molei Liu, Yi Zhang, and Doudou Zhou. Double/Debiased Machine Learning for Logistic Partially Linear Model, September 2020. URL `http://arxiv.org/abs/2009.14461`. arXiv:2009.14461 [stat].

[40] Lars van der Laan and Peter Gilbert. Semiparametric logistic regression for inference on relative vaccine efficacy in case-only studies with informative missingness. arXiv preprint arXiv:2303.11462, 2025. URL `https://arxiv.org/abs/2303.11462`.

[41] U Grenander. *Abstract Inference*. Wiley, New York. Wiley Series in Probability and Mathematical Statistics, 1981.

[42] Peter J. Bickel, Chris A.J. Klaassen, Ya'acov Ritov, and Jon A. Wellner. *Efficient and adaptive estimation for semiparametric models*. Johns Hopkins University Press Baltimore, 1993.

[43] JM Robins, A Rotnitzky, and LP Zhao. Estimation of regression-coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89:846–866, 1994.

[44] MJ Van der Laan and JM Robins. *Unified Methods for Censored Longitudinal Data and Causality*. Springer, 2003.

[45] Mark J van der Laan and Sherri Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York, 2011.

[46] Ted Westling, Mark J. van der Laan, and Marco Carone. Correcting an estimator of a multivariate monotone function with isotonic regression. *Electronic Journal of Statistics*, 14(2):3032–3069, January 2020. ISSN 1935-7524, 1935-7524. doi: 10.1214/20-EJS1740. URL https://projecteuclid.org/journals/electronic-journal-of-statistics/volume-14/issue-2/Correcting-an-estimator-of-a-multivariate-monotone-function-with-isotonic/10.1214/20-EJS1740.full.

[47] RL Prentice, JD Kalbfleisch, AV Peterson Jr, N Flournoy, et al. The analysis of failure times in the presence of competing risks. *Biometrics*, 34(4):541–554, December 1978.

[48] Shruti Khare, Céline Gurry, Lucas Freitas, Mark B Schultz, et al. GISAID's Role in Pandemic Response. *China CDC Weekly*, 3(49):1049–1051, 2021. ISSN 2096-7071. doi: 10.46234/ccdcw2021.255. URL http://weekly.chinacdc.cn/en/article/doi/10.46234/ccdcw2021.255.

[49] Torben Martinussen, Stijn Vansteelandt, and Per Kragh Andersen. Subtleties in the interpretation of hazard contrasts. *Lifetime Data Anal.*, 26(4):833–855, October 2020.

[50] Michael P Fay and Fan Li. Causal interpretation of the hazard ratio in randomized clinical trials. *Clinical Trials*, 21(5):623–635, October 2024. ISSN 1740-7745, 1740-7753. doi: 10.1177/17407745241243308. URL https://journals.sagepub.com/doi/10.1177/17407745241243308.

[51] Dean Follmann, Lauren Dang, and Michael Fay. Nonparametric identification of causal time-varying hazard ratios for vaccine field trials. Manuscript in preparation, 2026.

[52] Peter B. Gilbert, David C. Montefiori, Adrian B. McDermott, Youyi Fong, et al. Immune correlates analysis of the mRNA-1273 COVID-19 vaccine efficacy clinical trial. *Science*, 375 (6576):43–50, January 2022. doi: 10.1126/science.abm3425. URL https://www.science.org/doi/10.1126/science.abm3425.

[53] Craig A. Magaret, Li Li, Allan C. deCamp, Morgane Rolland, et al. Quantifying how single dose Ad26.COV2.S vaccine efficacy depends on Spike sequence features. *Nature Communications*, 15(1):2175, March 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-46536-w. URL https://www.nature.com/articles/s41467-024-46536-w.

[54] Fei Heng, Yanqing Sun, Li Li, and Peter B. Gilbert. Estimation and Hypothesis Testing of Strain-Specific Vaccine Efficacy With Missing Strain Types With Application to a COVID-19 Vaccine Trial. *Statistics in Medicine*, 44(6):e10345, March 2025. ISSN 0277-6715, 1097-0258. doi: 10.1002/sim.10345. URL https://onlinelibrary.wiley.com/doi/10.1002/sim.10345.

[55] Gellért Perényi and Mats Stensrud. Variant specific treatment effects with applications in vaccine studies. *Biometrics*, 81(2):ujaf068, April 2025. ISSN 0006-341X, 1541-0420. doi: 10.1093/biomtc/ujaf068. URL https://academic.oup.com/biometrics/article/doi/10.1093/biomtc/ujaf068/8155183.

[56] Leah I B Andrews, M Elizabeth Halloran, Kathleen M Neuzil, Lars van der Laan, Yunda Huang, Jessica Andriesen, Mayur Patel, Leigh H Fisher, Holly Janes, Nadine Rouphael, Stephen R Walsh, Deborah A Theodore, Hong-Van Tieu, Magdalena Sobieszczyk, Hana M El Sahly, Lindsey R Baden, Ann R Falsey, Thomas B Campbell, Colleen F Kelley, Catherine Mary Healy, Lilly Immergluck, Benjamin Luft, Ian Hirsch, Guy de Bruyn, Carla Truyers, Frances Priddy, Kelsey M Sumner, Brendan Flannery, Dean Follmann, Peter B Gilbert, and COVID-19 Prevention Network (CoVPN). Evaluating the test-negative design for COVID-19 vaccine effectiveness using randomized trial data: A secondary cross-protocol analysis of 5 randomized clinical trials. *JAMA Network Open*, 8(5):e2512763, May 2025.

[57] Kendrick Qijun Li, Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. Double Negative Control Inference in Test-Negative Design Studies of Vaccine Effectiveness. *Journal of the American Statistical Association*, pages 1–12, July 2023. ISSN 0162-1459, 1537-274X. doi: 10.1080/01621459.2023.2220935. URL https://www.tandfonline.com/doi/full/10.1080/01621459.2023.2220935.

[58] Torben Martinussen and Stijn Vansteelandt. On collapsibility and confounding bias in cox and aalen regression models. *Lifetime Data Anal.*, 19(3):279–296, July 2013.

Chapter 4

ADJUSTING FOR NEGATIVE CONTROL IMMUNE MARKERS TO ENHANCE EFFICIENCY OF EARLY-PHASE VACCINE TRIALS IN HIV-EXPOSED UNINFECTED INFANTS

Summary: *Adjustment for prognostic baseline variables can reduce bias due to covariate imbalance and increase efficiency in randomized trials. While the use of covariate adjustment in late-phase trials is justified by favorable large-sample properties, it is seldom used in small, early-phase studies, due to uncertainty in which variables are prognostic and the potential for precision loss, type I error rate inflation, and undercoverage of confidence intervals. To address this problem, we consider adjustment for a valid negative control outcome (NCO), or an auxiliary post-randomization outcome believed completely unaffected by treatment but correlated with the primary outcome. We articulate the assumptions that permit adjustment for NCOs without incurring post-randomization selection bias, and describe plausible data generating models where NCO adjustment can improve upon adjustment for baseline covariates alone. In numerical experiments, we illustrate performance and provide practical recommendations regarding model selection and finite-sample variance corrections. We apply our methods to the reanalysis of two early-phase vaccine trials in HIV exposed uninfected (HEU) infants, where we demonstrate that adjustment for auxiliary post-baseline immunological parameters can enhance precision of vaccine effect estimates relative to standard approaches that avoid adjustment or adjust for baseline covariates alone.*

4.1 Introduction

Vertical transmission of HIV-1 remains a global challenge in efforts to slow the HIV pandemic. Despite over a decade of use of highly effective antiretrovirals (ARVs) to prevent vertical transmission, 140,000 infant infections still occur annually worldwide,[1] with more than a third due to breastmilk transmission.[2,3] Identifying a safe, effective HIV-1 vaccine that can be administered to infants after birth to prevent transmission via breastmilk would be a major step towards eliminating perinatal HIV. Further, "germline targeting" HIV vaccines that induce broadly neutralizing antibodies (bnAbs) are under investigation, and recent work suggests that infants develop bnAbs more readily

than adults.[4–6] Comparing immune responses in infants and adults is informative for advancing both infant and adult HIV vaccine programs.

A key step towards identifying an effective HIV vaccine is evaluating its ability to induce cellular and humoral immune responses to the HIV-1 virus (hereafter referred to as a vaccine's *immunogenicity*). Evaluating vaccine immunogenicity in infants presents unique challenges. Unlike HIV-1 naïve adults, in whom humoral immune responses to HIV-1 under placebo can be assumed to be nearly zero,[7] HIV-exposed infants passively acquire HIV-specific antibodies from their mothers through the placenta. Therefore, the vaccine-elicited antibody response differs from the observed antibody response due to contamination by maternal antibodies. One approach to address this challenge is to measure antibody levels long enough after birth (roughly 12 months in most infants)[8] such that maternally inherited antibodies have waned. However, this approach limits evaluation of the dynamics of vaccine-induced antibody responses which may be useful for informing the optimal number and/or timing of vaccine doses. Another approach is to estimate the average vaccine effect on an immune marker of interest between infants randomly assigned to vaccine or placebo arms. However, small trial sizes, imbalanced randomization ratios used to maximize data on vaccine safety, and variability of antibody responses in both vaccine and placebo arms can yield imprecise vaccine effect estimates and low statistical power. More precise vaccine effect estimates could improve vaccine evaluations in HIV-exposed infants and other populations with prior pathogen exposure.

Covariate adjustment can potentially improve the precision of treatment effect estimates in randomized trials. Covariate-adjusted estimators of treatment effects can achieve asymptotic precision gains, supporting their application to large, late-phase trials.[9,10] Covariate adjustment may be especially beneficial in early-phase trials where sample sizes are constrained, chance covariate imbalance is more likely, and power to detect treatment effects may be low. However, covariate adjustment is seldom used in early-phase trials, since adjusting for weakly prognostic covariates can lead to precision loss, and adjustment for too many predictors can result in Type I error rate inflation and undercoverage of confidence intervals.[11–13] Current regulatory guidance recommends adjustment for a small number of covariates believed to be strongly associated with the outcome of interest.[14] Pooled analyses of several trials in adults have struggled to identify reliable baseline predictors of HIV-1 vaccine-elicited immune responses; Tte most prognostic baseline predictors (sex, age, and body mass index) exhibit at-best modest correlations with HIV-1-specific immune responses,[15,16]

and the applicability of these findings to infant studies remains unclear.

Auxiliary immune responses have the potential to more reliably predict vaccine-specific immune responses than commonly collected baseline variables. In the context of immune correlates analyses, Follmann[7] proposed imputing missing vaccine-elicited immune responses for placebo recipients using immune responses to an unrelated "baseline irrelevant vaccination". Positive correlation between immune responses to HIV-1 vaccines and tetanus toxoid and hepatitis B vaccinations have been observed in adults.[17] In populations with prior exposure to the pathogen of interest, baseline immune responses may be correlated with vaccine-elicited responses. Immune responses to herpes and influenza antigens at baseline were prognostic for immune responses at 6 weeks and 30 days after vaccination respectively.[18,19] Vaccine studies in HIV-exposed uninfected (HEU) infants can collect a variety of auxiliary immunological measurements. Baseline HIV-1-specific immune responses in both the infant and mother can serve as proxies for the level of passively-inherited antibody at birth. Immune responses to "off-target" HIV-1 antigens not included in the vaccine can describe the decay in maternal antibody over time. Immune responses to unrelated routine childhood vaccinations (e.g., tetanus, polio, haemophilus influenza B, hepatitis B, or measles) may be useful proxies for an infant's underlying immune function, or the potential to mount immune responses to the HIV-1 vaccine.

In an early-phase vaccine study, an analyst may be tempted to adjust for a small set of predictive auxiliary immunological measurements instead of baseline covariates to improve precision of the estimated vaccine effect. However, adjusting for variables measured post-randomization in a randomized trial is typically discouraged due to the potential of inducing post-treatment selection bias in treatment effect estimates.[20] We argue that adjustment for post-baseline variables can be accomplished while avoiding selection bias if the variable is completely unaffected by treatment. Borrowing terminology from the causal inference literature, we refer to these post-baseline variables unaffected by treatment as negative control outcomes (NCOs).[21,22] We focus on early-phase trials where sample sizes are small, data on the control condition is constrained by unequal randomization ratios, and precision gains are most sought. Our work offers the following contributions: (1) clear articulation of the assumptions enabling adjustment of post-baseline NCO, (2) demonstration of the asymptotic normality and semiparametric efficiency of the estimator that adjusts for a NCO, (3) causal graphical models describing instances where NCO adjustment can outperform baseline

covariate adjustment, (4) a simulation study examining performance of NCO adjustment and providing guidance on model selection and finite-sample variance corrections, and (5) comparison of our method to existing methods in the reanalysis of two early-phase HIV vaccine trials in HEU infants.

4.2 Methods

4.2.1 Potential Outcomes Framework

Consider a trial which randomly assigns a binary treatment A and measures a continuous primary outcome Y. Our arguments can be extended to cases where A takes k distinct levels or where Y is binary or count-based.[23] In addition to the primary outcome, suppose that the study measures baseline covariates X and an auxiliary outcome N. In our application, A represents an HIV-1 vaccine, and Y could be an antibody response to a HIV-1-specific antigen included in the vaccine construct. X could include baseline variables like infant birth weight and sex, while N could be an immune response to a routine childhood vaccine, an antibody response to a HIV-1-specific antigen not included in the vaccine, or a baseline immune response. Let $(Y(0), Y(1))$ and $(N(0), N(1))$ denote the potential outcomes for the primary and auxiliary outcomes under hypothetical placebo and vaccination.

With a continuous outcome and binary treatment, a typical causal estimand of interest is the population average treatment effect (ATE).

$$\text{ATE} = \mathbb{E}_0[Y(1) - Y(0)]$$

Suppose we adopt the following assumptions.

Assumption 12 (Stable Unit Transform Value Assumption (SUTVA).[24])**.** *SUTVA implies (i) no interference between participants, meaning that the potential outcomes of one participant does not depend on the treatment status of another, and (ii) no multiple versions of treatment.*

In some studies which measure infection or disease endpoints, the assumption of no interference between study participants is often questioned because an individual's infection outcome may

depend on the vaccination statuses of their close contacts.[25] However, our application focuses on immune response endpoints, which can safely be considered to be independent. No multiple versions of treatment is satisfied in double-blinded randomized trials with a sufficiently well-defined intervention. Under SUTVA, the observed data can be linked to the potential outcomes by the notion of consistency of potential outcomes; $Y = Y(A)$.

Assumption 13 (Known treatment propensity.[26])**.**

$$\pi = P(A = 1) \textit{ is known and lies between 0 and 1}$$

Assumption 13 implies that the treatment assignment is non-deterministic for all participants and is known by design. This is satisfied in the vast majority of randomized trials.

Assumption 14 (Strong ignorability.[27])**.**

$$A \perp (Y(0), Y(1), N(0), N(1), X)$$

Randomization ensures that treatment is independent of potential outcomes, covariates, and even unmeasured characteristics. Blinding participants, caregivers, and personnel involved in data collection and analysis prevents latent sources of post-randomization bias through participant behaviors, outcome ascertainment by caregivers, and participant/caregiver expectation of treatment benefit. Herein, we focus on randomized, double-blinded trials where the potential for bias due to incomplete blinding is unlikely.

The above assumptions are considered satisfied in well-designed and conducted randomized trials and serve as the basis for identification of the population ATE from randomized trial data. The following assumptions are not required for identification but can be leveraged to improve the precision of inference.

Assumption 15 (Negative control outcome (NCO).[20,22])**.**

$$N = N(0) = N(1)$$

Assumption 15 is strong and not typically invoked in randomized trials, because it implies

zero treatment effect on the auxiliary outcome in every participant. Because N is unaffected by treatment assignment, it can be effectively considered as a baseline covariate in subsequent developments.

Assumption 16 (No affect of Y on N)**.** *Consider the following nonparametric structural equations models (NPSEMs), which generated potential outcomes by a series of deterministic structural equations that depend on observed variables (X, A), possibly multivariate latent variables (U_N, U_Y), and random errors ($\epsilon_N \perp \epsilon_Y$).*

$$\begin{aligned} N &= f_N(X, U_N, \epsilon_N) \\ Y(a) &= f_Y(X, A = a, N, U_Y, \epsilon_Y) \end{aligned} \tag{4.1}$$

Where f_N, f_Y are unknown deterministic functions. We highlight three aspects of the NPSEMs. First, the NPSEMs are compatible with assumption 15 because the structural equation for N does not depend on the treatment variable A. Second, the structural equation for N does not depend on the realization of the primary outcome Y. If the structural equation for N depended on Y, assumption 15 would be violated because Y is affected by treatment. In the context of infant vaccine studies, N may represent a baseline immune response or a proxy for maternally inherited antibodies, while Y is a vaccine-elicited antibody response. In this case, N is unlikely to be affected by Y. If N is an immune response to an unrelated vaccine or an antigen not included in the vaccine, interference between antibody responses is possible, although evidence supports that infants and adults can mount immune responses to multiple distinct antigens simultaneously.[28] Third, we allow the structural equation for $Y(a)$ to depend on the observed values of N. Suppose N is a proxy for maternally inherited antibody. Previous research has shown that the presence of maternal inherited antibodies can "blunt" vaccine-elicited immune responses.[29] However, because maternally inherited antibodies are unaffected by vaccination, N may affect Y without introducing bias. In Subsection 4.2.3, we use causal directed acyclic graphs (DAGs) synonymous with the NPSEMs in Equation 4.1 to motivate the potential benefits of adjusting for N in a randomized experiment.

4.2.2 Baseline Covariate Adjustment

Consider a randomized trial with binary treatment A, baseline covariate X, and primary outcome Y. We assume the observed data are n independent and identically distributed (i.i.d.) samples from some distribution P_0 in statistical model M.

$$O_i = (A_i, X_i, Y_i) \overset{iid}{\sim} P_0 \in M$$

A natural estimator for the population ATE is the plug-in estimator, or the difference in means between the two treatment groups.

$$\hat{\psi}_{\text{plug-in}} = \bar{Y}(1) - \bar{Y}(0) = \frac{1}{n}\sum_{i=1}^{n} \frac{A_i Y_i}{\hat{\pi}} - \frac{(1-A_i)Y_i}{(1-\hat{\pi})} \tag{4.2}$$

Where $\bar{Y}(0)$ and $\bar{Y}(1)$ are the observed mean outcomes in the placebo and treatment group respectively, and $\hat{\pi} := \frac{1}{n}\sum_{i=1}^{n} A_i$ is the observed treatment probability. Tsiatis et al.[9] demonstrated that any reasonable estimator of the population average treatment effect in a randomized trial measuring baseline covariates X either belongs to or is asymptotically equivalent to a member of the following class of augmented inverse probability weighted (AIPW) estimators.

$$\hat{\psi}_{\text{AIPW}} \in \left\{ \hat{\psi}_{\text{Plug-in}} - \frac{1}{n}\sum_{i=1}^{n} (A_i - \hat{\pi}) \left\{ \frac{h_0(X_i)}{1-\hat{\pi}} + \frac{h_1(X_i)}{\hat{\pi}} \right\} : \text{Var}(h_a) < \infty \right\} \tag{4.3}$$

Where h_0 and h_1 are arbitrary functions with finite variance. The AIPW estimators modify the plug-in estimator by subtracting a mean-zero augmentation term that exploits baseline covariates to improve precision. Tsiatis et al.[9] derived the smallest possible variance achievable by an estimator in the class in 4.3. The *semiparametric efficient estimator* is the estimator which achieves the minimum variance bound, and is given by $\hat{\psi}_{\text{AIPW}}$ with $h_a = \mathbb{E}[Y|A=a, X]$. The derivation of the variance bound and efficient estimator are based on arguments rooted in semiparametric theory, a review of these concepts and the derivation can be found in Appendix C.2.

In practice, one must propose models for the unknown functions $h_a(X)$. In the case of a continuous outcome where the number of covariates is small, a common approach uses ordinary least squares (OLS) regressions for $\hat{h}_a(X)$, fit separately in each arm or equivalently with full treatment-

covariate interactions.[9,10] The resulting estimator $\hat{\psi}_{\text{Cov-AIPW}}$ is consistent and asymptotically normal even if the models for $h_a(X)$ are misspecified, because the estimator belongs in the class of estimators in 4.3. Hence, $h_a(X)$ are often referred to as *working models*, as correct specification is not required to attain consistency and normality. When OLS-working models are used, $\hat{\psi}_{\text{Cov-AIPW}}$ achieves guaranteed efficiency gain relative to plug-in estimators in large samples.[10] Alternatively, nonlinear regression models (e.g., GLMs or cross-fitted machine learning estimators) could be used as working models for adjustment, and can achieve guaranteed efficiency gain when paired with an OLS calibration procedure applied to the model predictions.[30,31] To obtain valid confidence intervals and hypothesis tests, a consistent estimator of the variance of $\hat{\psi}_{\text{Cov-AIPW}}$ is required. Motivated by our application to early-phase trials, we also desire a variance estimator that is robust to heteroskedasticity, misspecification of the working models, and exhibits good performance in finite samples. Tsiatis advocated for the use of sandwich standard errors with a degrees-of-freedom correction to account for the additional variability induced by model-fitting.[9] We consider the general variance estimator.[32]

$$\widehat{\text{Var}}(\hat{\psi}_{\text{Cov-AIPW}}) = (\mathbf{C} \odot \mathbf{R})^T \cdot (\mathbf{C} \odot \mathbf{R})$$

Where $\odot$ denotes element-wise product, $\mathbf{C} = (C_1, \ldots, C_n)^T$ is a vector of finite-sample correction factors, and $\mathbf{R} \in \mathbb{R}^n$ are the usual components of the sandwich variance with entries

$$\begin{aligned} R_i = &\left(\frac{A_i}{\sum_{j=1}^n A_j} - \frac{(1-A_i)}{\sum_{j=1}^n (1-A_j)} \right) Y_i - \frac{\hat{\psi}_{\text{AIPW}}}{n} - (A_i - \hat{\pi}) \left(\frac{\hat{h}_0(X_i)}{\sum_{j=1}^n (1-A_j)} + \frac{\hat{h}_1(X_i)}{\sum_{j=1}^n A_j} \right) \\ &- (A_i - \hat{\pi}) \left(\frac{\bar{Y}(0) - \bar{h}_0}{\sum_{j=1}^n (1-A_j)} + \frac{\bar{Y}(1) - \bar{h}_1}{\sum_{j=1}^n A_j} \right) \end{aligned}$$

Where $\bar{h}_0 := \frac{1}{n} \sum_{i=1}^n \hat{h}_0(X_i)$, $\bar{h}_1 := \frac{1}{n} \sum_{i=1}^n \hat{h}_1(X_i)$. While variance corrections become negligible in large samples, they may be helpful in avoiding undercoverage of confidence intervals and type I error rate inflation in finite samples.[11] To improve finite-sample performance, we explore the following correction factors.

1. HC0-type: $C_i = 1$, no correction.

2. HC1-type: $C_i = \left(\frac{(n_0-p_0-1)^{-1}+(n_1-p_1-1)^{-1}}{(n_0-1)^{-1}+(n_1-1)^{-1}}\right)^{1/2}$ where p_0 and p_1 are the number of parameters fit in models $\hat{h}_0$ and $\hat{h}_1$ respectively, and n_1, n_0 are the number of treated and control participants respectively.[9]

3. HC2-type: $C_i := (1/(1 - H_{a,i,i}))^{1/2}$ where $H_{a,i,i}$ is the leverage of observation i, or the i-th diagonal entry of the "hat matrix" for the OLS regression in arm $A = a$.[32]

4. HC3-type: $C_i := 1/(1 - H_{a,i,i})$ where $H_{a,i,i}$ is defined identically above.[32]

4.2.3 Negative Control Outcome Adjustment

Covariate adjustment in randomized experiments almost exclusively focuses on pretreatment variables. Adjusting for a post-treatment variable introduces the possibility that the variable was affected by treatment, which can result in selection bias. However, fixating on pretreatment variables is sufficient but not necessary to avoid bias. The hallmark of acceptable adjustment variables is not that they were measured before the intervention, but that they were unaffected by the intervention.[20] We formalize this notion using the NCO Assumption (Assumption 4). NCOs have been used extensively for bias detection and elimination in non-randomized studies,[21,22] but we consider how NCOs can be used to augment inference in randomized trials.

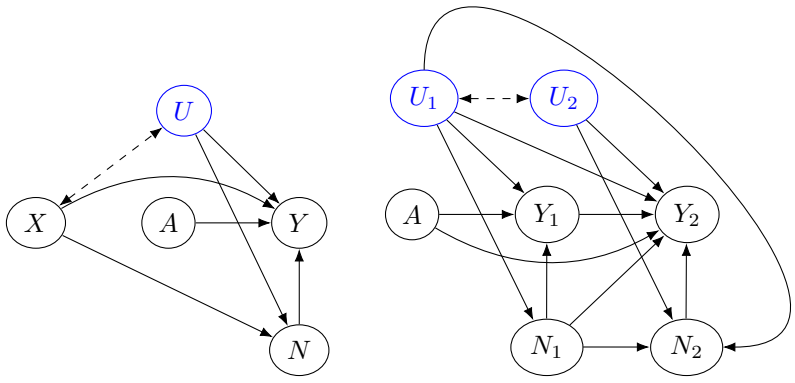

Figure 4.1: Causal diagrams illustrating assumptions where NCOs are a proxy for a prognostic unmeasured precision variable U (left) or prognostic unmeasured time-varying process (U_1, U_2) (right). In the right panel, baseline covariates X are omitted for simplicity.

We provide two examples of graphical causal models based on directed acyclic graphs (DAGs) where NCO adjustment may offer benefits relative to baseline covariate adjustment. First, we focus on the DAG in the left panel of Figure 5.1. The DAG corresponds to the NPSEMs in Equation 4.1 when the unmeasured variables affecting the primary and negative control outcomes are identical, i.e., $U = U_Y = U_N$. In the literature on bias detection in observational studies using negative controls, the degree of overlap between U_Y and U_N is often referred to as the "U-comparability" of N and Y.[22] When N and Y are exactly U-comparable and the contributions of the errors to each outcome are small, N is a reliable surrogate for the common unobserved variable U affecting the primary outcome.[20] In reality, N and Y may only be approximately U-comparable (i.e., $U_N \approx U_Y$), but N may still be a useful proxy for the effect of U_Y on Y. In our application to infant HIV vaccine studies, an important missing variable (U) is the function of an infant's immune system at birth. While baseline covariates such as infant sex or birth weight will capture limited information on U, immune responses to routine childhood vaccinations (N) are unlikely to be affected by HIV-1 vaccination but may be relevant proxies for U. A second causal model for longitudinal data is described in the DAG in the right panel of Figure 5.1. In the DAG, N and Y are exactly U-comparable and the tuple (N_1, N_2) can be considered as a proxy for the time-varying unmeasured process (U_1, U_2). In infant vaccine studies, the decay of maternally inherited antibodies is a plausible example of a time-varying unmeasured process $(U_1, U_2, \ldots)$. Longitudinal measures of immune responses to synthetic antigens or HIV-1 antigens not included in the vaccine $(N_1, N_2, \ldots)$ may be useful surrogates for maternal antibody decay. Covariates measured at baseline may be particularly poorly equipped to explain how immune responses will evolve over time as maternal antibody decays.

We highlight a few reasons why adjustment for NCOs may be appealing from an analyst's perspective seeking to reduce finite-sample bias and improve precision of ATE estimates. First, an NCO measured during follow-up may be more predictive of the outcome of interest than baseline variables, especially when participant characteristics evolve post-baseline (e.g., due to developmental or maturation processes). If the NCO is a proxy for such changes, it may be a better choice of adjustment variable than baseline predictors. Second, the NCO may explain idiosyncratic, technical sources of variation in the primary outcome that cannot be explained by baseline variables. In our application to vaccine studies, binding antibody multiplex assays (BAMA) are often used to

measure antibody responses to several antigens in parallel. A negative control antibody response measured using BAMA may explain variability in the primary antibody responses due to sample preparation and assay execution. Third, NCOs can simplify the problem of selecting which variables to adjust for. Randomized trials commonly collect many baseline covariates with questionable relevance to the outcome of interest. While adjustment for all predictors will not harm large-sample efficiency (see Proposition 3 below), it can lead to efficiency loss in finite-samples when variables exhibit low or moderate correlation with the primary outcome. Thus, adjusting for a small, valid set of predictive NCOs may be preferable to adjusting for a large set of baseline covariates in a small trial.

We present the following Proposition, which extends the asymptotic theory for covariate adjustment to include NCOs. A proof of Proposition 3 can be located in Appendix C.2.2.

Proposition 3. *If Assumptions assumption 12 to 16 hold, the following model-assisted estimator is consistent and asymptotically normal for the ATE for any choices of working models $\hat{h}_0, \hat{h}_1$ with finite variance.*

$$\hat{\psi}_{CovNCO\text{-}AIPW} := \hat{\psi}_{Plug\text{-}in} - \frac{1}{n}\sum_{i=1}^{n}(A_i - \hat{\pi})\left\{\frac{\hat{h}_0(X_i, N_i)}{1-\hat{\pi}} + \frac{\hat{h}_1(X_i, N_i)}{\hat{\pi}}\right\} \tag{4.4}$$

If $\hat{h}_0$ and $\hat{h}_1$ are fit using OLS regression with treatment-predictor interactions, the estimator will have guaranteed asymptotic efficiency gain over unadjusted estimators and estimators which adjust for X alone. If $\hat{h}_0$ and $\hat{h}_0$ are correctly specified models for $\mathbb{E}[Y|X, N, A = 0]$ and $\mathbb{E}[Y|X, N, A = 0]$, $\hat{\psi}_{CovNCO\text{-}AIPW}$ will be semiparametric efficient, meaning it achieves the lowest asymptotic variance among all asymptotically linear estimators and data generating laws satisfying $A \perp (X, N)$.

Proposition 3 indicates the AIPW estimator that adjusts for (X, N) is the most efficient estimator of the ATE in large samples among trials that measure (X, N). However, we can show that adjustment for an NCO alone can lead to semiparametric efficient inference in an augmented data model under additional assumptions detailed in the corollary below.

Corollary 1. *Consider a model M^* encompassing data generating laws for a randomized trial measuring an oracle data unit that includes the unmeasured variable U: $\{X_i, N_i, U_i, A_i, Y_i\}_{i=1}^n \overset{i.i.d}{\sim} P_0 \sim M^*$. Suppose assumption 12 to 16, where ignorability with repsect to U, N also hold:*

$\{Y(1), Y(0), N, X, U\} \perp A$. *The following model-assisted estimator is consistent and asymptotically normal for the ATE for any choices of working models* $\hat{h}_0, \hat{h}_1$ *with finite variance.*

$$\hat{\psi}_{NCO\text{-}AIPW} \coloneqq \hat{\psi}_{Plug\text{-}in} - \frac{1}{n}\sum_{i=1}^{n}(A_i - \hat{\pi})\left\{\frac{\hat{h}_0(N_i)}{1-\hat{\pi}} + \frac{\hat{h}_1(N_i)}{\hat{\pi}}\right\} \tag{4.5}$$

When $\hat{h}_0, \hat{h}_1$ *are OLS regression models fit with full treatment-predictor interactions, then* $\hat{\psi}_{NCO\text{-}AIPW}$ *offers guaranteed efficiency gain relative to the plug-in estimator.*

In addition, if we assume (i) N *and* Y *are U-comparable,* $U = U_Y = U_N$ *in Equation 4.1, and (ii)* N *is a surrogate for the effect of* (X, U) *on* Y *within levels of* A*, implied by the following mean independence assumption.*

$$E[Y|A, X, U, N] = E[Y|A, N]$$

When $\hat{h}_0, \hat{h}_1$ *are OLS regression models fit separately within each treatment arm, then* $\hat{\psi}_{NCO\text{-}AIPW}$ *exhibits guaranteed efficiency gain relative to the unadjusted and covariate-adjusted estimator using OLS-working models and no efficiency loss relative to the OLS-assisted estimator which adjusts for the oracle* $\{X, U, N\}$*. Furthermore, if* $\hat{h}_0, \hat{h}_1$ *are correctly specified models for* $E[Y|N, A = 0]$ *and* $E[Y|N, A = 1]$*, then* $\hat{\psi}_{NCO\text{-}AIPW}$ *is efficient in the oracle model, meaning it achieves the lowest asymptotic variance among asymptotically linear estimators across all data generating laws for the oracle data unit* $\{A, X, U, N, Y\}$ *satisfying* $A \perp (X, U, N)$*.*

Corollary 1 highlights several important characteristics of the parsimonious estimator that only adjusts for the NCO. First, adjustment for a valid NCO will lead to improved efficiency relative to an unadjusted estimator. Second, if N is a valid surrogate for both measured and unmeasured causes of Y, then adjusting for the NCO improves upon adjusting for baseline covariates and will not lose efficiency relative the hypothetical oracle estimator that adjusts for all measured and unmeasured causes. Third, if the working models are specified correctly, adjustment for N alone leads to the most efficient estimator among all estimators adjusting for (X, U, N).

We acknowledge that the NCO assumption is a strong assumption not typically invoked when analyzing randomized trials. There exists a tension between choosing an auxiliary outcome that is sufficiently predictive of the primary outcome but that is not affected by the intervention. The

most predictive auxiliary outcomes may be more likely to introduce bias. The plausibility of the NCO assumption may hinge on subject matter knowledge and evidence from prior experiments, which may be limited in early-phase trials. However, the NCO assumption may be easier to justify in vaccine studies in specific circumstances. Empirical evidence from late-phase efficacy trials of HPV[33] and COVID-19 vaccines[34] have shown negligible vaccine effects in preventing infections caused by vaccine untargeted pathogens, supporting their use as NCOs. Previous HIV-1 vaccine studies in infants have shown that markers of maternal antibody[35,36] and immune responses to routine childhood vaccinations[28] are unaffected by HIV-1 vaccination. However, identifying valid NCOs may be more difficult in other circumstances. Many modern HIV-1 vaccines aim to elicit immune responses that confer broad protection against many HIV-1 strains. In such cases, immune responses to HIV-1 antigens not included in the vaccine may be risky choices of NCOs. Choosing an NCO immune marker sufficiently distinct from the vaccine's mechanism of action may be an easier task for vaccines with fewer components (e.g., mRNA/DNA, protein subunit, or conjugate vaccines) than more antigenically diverse constructs (e.g., live attenuated, inactivated, mosaic, or polyvalent). In some rare cases, vaccines can confer protection against pathogens unrelated to the target disease through a memory-like response developed in innate immune cells (e.g., BCG vaccination given to infants).[37] Ultimately, the selection of NCOs should be based on scientific knowledge and available data from prior studies.

If an analyst wants to check the NCO assumption prior to adjustment, we suggest employing a pretest. We restrict focus to tests conducted at the nominal Type I error rate ($\alpha = 0.05$). Let $\tau^N = N(1) - N(0)$ refer to the individual treatment effect of vaccine on the candidate NCO. One can test the sharp causal null hypothesis implied by Assumption 4 that $\tau^N = 0$ for all participants using a randomization test, which envisions the treatment assignment mechanism as the sole source of randomness in an experiment. Rejection of the sharp null indicates that the observed data are unlikely under the NCO assumption, and NCO adjustment should be avoided. However, failure to reject the sharp null cannot be interpreted as evidence that the NCO assumption is true, and small early-phase trials may possess low power to discriminate between valid and invalid NCOs. As an alternative, we consider an *equivalence pretest* of the hypothesis that individual causal effects τ^N lie outside an analyst-defined "equivalence window" centered at zero, $[-\epsilon, \epsilon]$ for some analyst-chosen threshold $\epsilon > 0$. Formally, we define the null hypothesis as $H_0^{\text{equiv}} : \tau^N < -\epsilon \cup \epsilon < \tau^N$. Rejection

of H_0^{equiv} implies that individual treatment effects are not too large, while failure to reject H_0^{equiv} implies that the analyst cannot rule out large individual treatment effects. Note that H_0^{equiv} is a hypothesis of *bounded* individual treatment effects, for which valid tests exist when the test statistic used for randomization inference satisfies certain regularity conditions.[38] Unlike the sharp null, the equivalence null hypothesis assumes that N is an *invalid* NCO, endowing the equivalence test with skepticism towards candidate NCOs in small studies when small equivalence thresholds ϵ are chosen. To guide choices of ϵ, we propose using fractions of measures of dispersion (e.g., range or variance) of the primary outcome pooled across treatment arms. In forthcoming numerical experiments, we will show the relative benefits and drawbacks of conducting pretests before NCO adjustment.

4.2.4 Finite-sample inference

The approaches developed above are based on a statistical model that assumes units are exchangeable draws from a hypothetical superpopulation, and the parameter of interest is a treatment effect defined in the superpopulation. In early-phase trials, the superpopulation sampling model may be implausible if the trial imposes strict eligibility criteria or if the participants represent a convenience sample. An analyst may wish to restrict focus to whether the intervention worked in the study participants themselves. We explore two different approaches to finite-population inference.

Adjusted Estimators of the Sample Average Treatment Effect

Adopting Neyman's randomization inference framework,[39] the potential outcomes $(Y(1), Y(0))$ and predictors (X, N) are assumed fixed, and treatment assignment mechanism is the sole source of randomness in the study. Our estimand of interest is the average treatment effect in the study participants themselves, or the sample average treatment effect (SATE).

$$\text{SATE} = \frac{1}{n}\sum_{i=1}^{n} Y_i(1) - Y_i(0) \tag{4.6}$$

One may be interested in estimating the SATE and testing the *weak null hypothesis*, H_0^{Weak} : SATE $= 0$.[39] The SATE is not identified because only one potential outcome is observed per individual. However, the plug-in estimator given in Equation 4.2 is unbiased for the SATE and a conservative variance estimator is $\frac{\hat{S}^2_{Y,A=1}}{n_1} + \frac{\hat{S}^2_{Y,A=0}}{n_0}$ where $\hat{S}^2_{Y,A=a}$ is the sample variance of the

outcome in arm a. Lin[40] proposed a family of linearly adjusted estimators, which we extend to include NCOs.

$$\hat{\psi}_{\text{Lin}} \in \left\{ \frac{1}{n_1} \sum_{i=1}^{n} A_i (Y_i - \beta_1^T (X_i - \bar{X}, N_i - \bar{N})) - \frac{1}{n_0} \sum_{i=1}^{n} (1 - A_i)(Y_i - \beta_0^T (X_i - \bar{X}, N_i - \bar{N})) \right\}$$

Where $\bar{X}$ and $\bar{N}$ are the sample means of the baseline covariates and NCO in the study population respectively. The choice of (β_0, β_1) that minimizes the variance is obtained by OLS regression of outcomes Y on predictors (X, N) separately within each treatment arm A, or equivalently with full treatment-predictor interactions. The resulting estimator is consistent, asymptotically normal, and yields efficiency gain relative to the plug-in estimator in large samples under bounded moment and asymptotically stable randomization probabilities.[40] Lin also advocated estimating the variance of $\hat{\psi}_{\text{Lin}}$ using the sandwich method, which is consistent, asymptotically conservative, and robust to heteroskedasticity and linear model misspecification.

There are some caveats to using Lin's estimator for estimating the SATE in finite samples. First, the adjusted estimator introduces finite-sample bias due to estimating (β_0, β_1). The bias is absent when the parameters are known. However, Lin showed that the bias diminishes rapidly with increasing sample size and argued that the bias due to adjustment should be weighed against the bias from covariate imbalance. Second, in small samples and/or settings with high leverage points due to outlying covariate values, sandwich standard errors may be anti-conservative. Finite-sample variance corrections, as discussed above, can restore the coverage of confidence intervals (CIs) and Type I error rate control at the nominal level. Third, Lin's sandwich variance estimator cannot be used for superpopulation inference because it does not account for the additional uncertainty due to centering the predictors. In randomization inference however, covariate centering can be ignored because the mean of the predictors is assumed fixed. Variance estimators which account for covariate centering for superpopulation inference have been developed elsewhere.[10]

Adjusted Randomization Tests

We continue within Neyman's causal model, where all covariates and outcomes are fixed and the only source of randomness in the experiment is the treatment assignment mechanism. Rather than estimating a treatment effect, suppose interest lies in testing whether the intervention satisfies the

sharp null hypothesis of a zero treatment effect in study participants.

$$H_0^{\text{sharp}} : Y_i(1) - Y_i(0) = 0 \quad \text{for all i} \in \{1, \ldots, n\}$$

Under H_0^{sharp}, both potential outcomes $\{Y_i(0), Y_i(1)\}$ are known for every participant. The exact distribution of any arbitrary test statistic $T : \{Y_i(0), Y_i(1)\}_{i=1}^n \to \mathbb{R}$ under H_0^{sharp} can be obtained by computing the test statistic over all possible permutations of the treatment assignment vector $\mathbf{A}$. If the total number of permutations of $\mathbf{A}$ is prohibitively large, the null distribution can be approximated using Monte Carlo methods. The test statistic computed on the observed data can be compared to randomization distribution under H_0^{sharp} to construct finite-sample exact tests of the sharp null hypothesis.

There are two main strategies for adjusting for predictors in randomization tests. We refer to the first approach as the pseudo-outcome approach. In this approach, an arbitrary algorithm f regresses the outcomes Y on non-treatment predictors (X, N) to produce residuals $r = Y - f(X, N)$. The residuals $\{r_i\}_{i=1}^n$ are then used as the basis for randomization inference in a covariate-free fashion. The pseudo-outcome approach can lead to improved power of the randomization test when the residuals are more stable and less dispersed than the outcomes.[41] Under the Neyman model, covariates, valid NCOs, and potential outcomes are presumed fixed. Therefore, functions of these variables, including the residuals computed from the regression fit, are also fixed.[41] The regression f is not a stochastic model but is merely an algorithmic fit to fixed quantities and introduces no additional uncertainty to the experiment. The pseudo-outcome approach can be used with any regression algorithm – linear regression,[42] nonlinear regression,[43] or data-adaptive methods[44,45] with or without variable selection techniques.

We refer to the second strategy as the model-output approach. The model-output approach involves regressing the outcome Y on predictors (X, N) and treatment A, then using the coefficient on the treatment variable A as the test statistic upon which to base randomization inference.[43,46] Compared to the pseudo-outcome approach which requires only one algorithmic fit, the model-output approach requires refitting a model to each dataset corresponding to permutations of the treatment assignment vector, and may therefore incur additional computational costs. In theory, any model could be used to summarize the treatment effect, but the canonical choice is the linear

model. As described by Zhao and Ding, summarizing the treatment effect using the robust t-statistic based on Lin's estimator of the SATE with OLS models and full interactions between treatment and predictors offers a number of theoretical advantages.[46] That is, we can base randomization inference on the following quantity

$$T = \frac{\hat{\psi}_{\text{Lin}}}{\sqrt{\hat{S}_0^2/n_0 + \hat{S}_1^2/n_1}}$$

Where $\hat{S}_a^2 := (n_a - 1)^{-1} \sum_{i:A_i=a} (Y_i - \hat{Y}(a))^2$ is the sample variance in arm $A = a$. No additional finite-sample corrections are required to ensure that randomization inference using Lin's robust t-statistic provides an exact test for H_0^{sharp} in finite samples. The procedure is also asymptotically valid for testing the weak null hypothesis of zero *average* treatment effect in the sample (H_0^{weak} : SATE = 0), which significantly broadens the test's application beyond the sharp null. Third, the procedure is guaranteed to be more powerful than unadjusted randomization tests under all alternatives hypotheses even if the OLS models are misspecified.[46] In subsequent numerical experiments, we will focus on the model-output strategy based on Lin's t-statistic, given its desirable theoretical properties.

4.3 Numerical Experiments

We simulated trials ranging from small to moderate in size: $n \in \{40, 60, 80, 100, 120\}$ and considered a randomization probability favoring the intervention ($\pi = 0.8$). For simplicity, we assumed a binary treatment A with constant additive treatment effect parameterized by $\beta = 1$. Our simulated trials measured a single quantitative baseline covariate X, primary outcome Y, and an auxiliary outcome N at a single point in time. We focus on estimating the population average treatment effect; results of numerical experiments for finite-sample inference can be found in the Appendix.

Each observation was generated as i.i.d. draws from the following model.

$$(X, U) \sim N\left(\mu = (0, 0), \Sigma = \begin{pmatrix} 1 & \rho_{(X,U)} \\ \rho_{(X,U)} & 1 \end{pmatrix}\right)$$

$$g(N) \sim \text{Norm}(\mu_N = \beta_0 + \beta_1^N X + \beta_2^N U, \sigma^2 = 1)$$

$$Y(0) \sim \text{Norm}(\mu_Y = \beta_0 + \beta_1^Y X + \beta_2^Y U, \sigma^2 = 1)$$
$$Y(A) = Y(0) + \beta A$$

For the main simulation, the auxiliary outcome N was assumed to be a valid NCO satisfying Assumptions 4 and 5. We consider cases where g is either the identity function (Setting 1) or a logistic function, $g(x) = -\log(8/x - 1)$ (Setting 2). The logistic function implies a distribution saturated at both lower and upper detection limits, a scenario common in immunogenicity studies. We assumed $\rho_{(X,U)} = 0$ implying $X \perp U$ and assumed $\beta_0 = 1$ was fixed. We also assumed that $\beta_1^N = \beta_1^Y$ and $\beta_2^N = \beta_2^Y$, implying that the measured and unmeasured variables affected N and $Y(0)$ with equal magnitude. This is similar to additive equiconfounding assumptions invoked for bias adjustment using NCOs in observational studies.[22] Using properties of the multivariate normal distribution, we express the remaining data-generating parameters, β_1 and β_2, in terms of more interpretable quantities (see Appendix C.1.3 for details).

$$\beta_1 = \sqrt{\frac{\rho^2_{(Y,X)}(\beta_2^2 + 1)}{1 - (\rho_{(Y,X)})^2}} \quad \beta_2 = \sqrt{\frac{\rho_{(Y,N|X)}}{1 - \rho_{(Y,N|X)}}}$$

Where $\rho_{(Y,X)}$ describes the correlation between Y and X and $\rho_{(Y,N|X)}$ captures the correlation between the NCO and placebo outcome over and above that explained by the measured covariate. The baseline covariate was weakly predictive ($\rho_{(Y,X)} = 0.3$) and we varied $\rho_{(Y,N|X)} \in \{0, 0.3, 0.5, 0.8\}$ encompassing cases where the N was not predictive, weakly predictive, moderately predictive, and strongly predictive of $Y(0)$.

We compared the performance of the following estimators: (a) Plug-in: $\hat{\psi}_{\text{plug-in}}$; (b) Covariate-adjusted: $\hat{\psi}_{\text{AIPW}}$ with $\hat{h}_a(X)$ fit using OLS regression with covariate-treatment interactions; (c) NCO-adjusted: $\hat{\psi}_{\text{AIPW}}$ with $\hat{h}_a(N)$ fit using OLS regression with predictor-treatment interactions; (d) Quantile-NCO-adjusted: $\hat{\psi}_{\text{AIPW}}$ with $\hat{h}_a(N^*)$ fit using OLS regression with predictor-treatment interactions after the empirical quantile transform, $N^* = \hat{F}_N(N)$, was applied to the NCO; and (e) Fully adjusted: $\hat{\psi}_{\text{AIPW}}$ with $\hat{h}_a(X, N)$ fit using OLS regression with predictor-treatment interactions. Estimators were compared with respect to (i) absolute finite-sample bias as a function of sample size n and $\rho_{Y,N|X}$ relative to the plug-in estimator; (ii) average confidence intervals coverage

over varying sample sizes n and finite sample corrections C; (iii) relative efficiency defined as the ratio of each estimator's estimated variance compared to the plug-in estimator; and (iv) Power and Type I error of Wald tests of the null hypothesis that the population average treatment effect was zero under varying effect sizes and $\rho_{Y,N|X}$. We summarized simulation results over 1000 simulation replicates per condition.

The results of the numerical experiments are shown in Figure 4.2. In the left panels of Figure 4.2, NCO-adjusted estimators incurred greater finite-sample bias and lower precision in small trials when the NCO was skewed and not prognostic. Adjustment for both covariates and weakly prognostic NCOs generated additional bias and precision loss relative to simpler adjustment approaches. However, NCO-adjusted estimators achieved reductions in finite-sample bias and precision enhancement when sample sizes were moderate and the NCO was moderately correlated with the primary outcome. Adjustment for a quantile-transformed NCOs led to the largest finite-sample bias reduction and largest precision gain when the NCO was skewed. When the NCO was not skewed, adjustment for a quantile-transformed NCO still achieved meaningful improvements in performance. In the upper right panel of Figure 4.2, the HC3 correction produced coverage estimates that matched or exceeded that of the plug-in estimator across all estimators, sample sizes, and data-generating mechanisms, supporting its use in early-phase trials. Lastly, in the bottom right panel, we observed that NCO adjustment can lead to improvements in power to test $H_0 : \text{ATE} = 0$ for $n = 60$ even when the NCO was weakly or moderately predictive of the outcome of interest. Adjusting for a quantile-transformed NCO achieved the highest power when the NCO distribution was skewed and achieved power rivaling adjustment for an untransformed NCO when the NCO distribution was not skewed. Results for estimating the SATE and for randomization inference were very similar and can be located in the Appendix C.1.

We also conducted numerical experiments examining the performance of NCO adjustment when the NCO assumption (Assumption 15) was violated. Full details and results can be found in the Supplementary materials. We considered pairing the NCO-adjusted estimator with either a pretest of the sharp null hypothesis or an equivalence pretest with different equivalence thresholds $\epsilon > 0$. All pretests were conducted at level $\alpha = 0.05$. If the sharp null hypothesis was rejected, the plug-in estimator would be returned, otherwise the NCO-adjusted estimator was returned. If the equivalence null hypothesis was rejected, the NCO-adjusted estimator was returned and otherwise,

the plug-in estimator was returned. As anticipated, NCO adjustment resulted in bias when the treatment affected the candidate NCO. Checking the NCO assumption using a pretest partially mitigated the bias. A pretest of the sharp null protected against bias and poor coverage, especially when the NCO exhibited low correlations with the primary outcome. Equivalence pretests protected against violations of the NCO assumption outside the equivalence window $[-\epsilon, \epsilon]$. When NCO violations occurred within the equivalence window, bias occurred particularly for moderate size trials. When the NCO assumption was correct, pretests dampened the benefits of NCO adjustment. Pairing NCO adjustment with a pretest of the sharp null or an equivalence pretest (using a wide equivalence threshold) resulted in intermediate efficiency gains between the plug-in estimator and the NCO-adjusted estimator.

The implications of our simulation study for early-phase randomized trials are (1) NCO adjustment can lead to improved precision and power in small samples when the NCO is valid and is predictive of the primary outcome, (2) in the presence of skewed predictors subject to detection limits, adjustment for a quantile transformed NCO can mitigate finite-sample bias and variance inflation, (3) applying the HC3 correction to sandwich standard errors provided robust inference in finite-sample settings, and (4) despite the large-sample optimality of fully-adjusted estimators (as described in Proposition 3), more parsimonious adjustment strategies offer better performance than fully-adjusted estimators in finite samples.

4.4 Application to two early-phase vaccine trials in HIV-Exposed Uninfected Infant

We apply our proposed methodology to two early-phase vaccine trials involving HIV exposed uninfected (HEU) infants.

4.4.1 HPTN 027

The HIV Prevention Trials Network (HPTN) 027 was a randomized, double blind, placebo-controlled Phase I safety and immunogenicity study evaluating ALVAC-HIV vCP1521 (ALVAC) versus placebo in 60 HEU infants at Mulago National Referral Hospital in Kampala, Uganda.[47] Notably, HPTN 027 was the first perinatal HIV vaccine trial conducted in Africa and evaluated the ALVAC vaccine used in the Thai RV144 trial, which demonstrated 30% protection against HIV-1 acquisition in

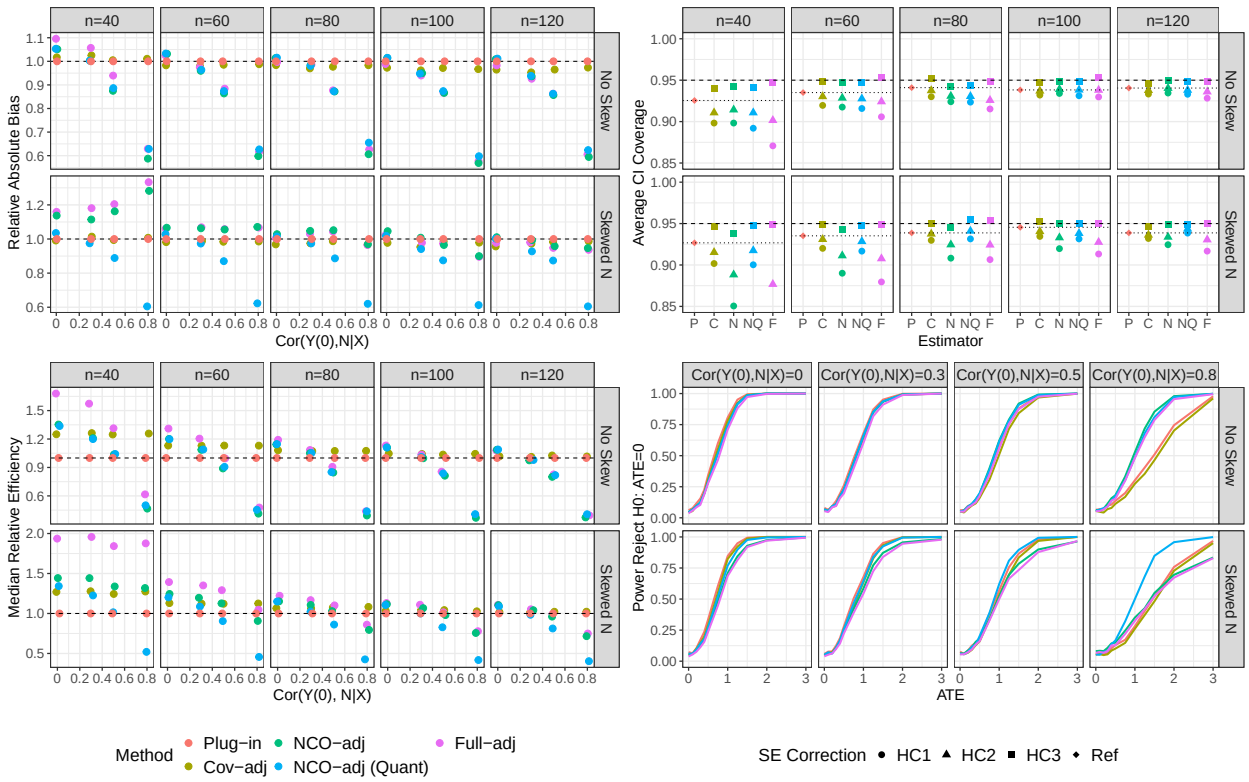

Figure 4.2: Top left: relative absolute bias of different estimators as a function of sample size, data generating mechanism, and predictiveness of NCO. Top right: average coverage of 95% confidence intervals for different estimators with different finite-sample variance corrections as a function of sample size and data generating mechanism. Bottom left: median relative efficiency of different estimators as a function of sample size, data generating mechanism, and predictiveness of NCO. Bottom right: power curves of Wald tests of H_0 : ATE $= 0$ as a function of NCO predictiveness and the data generating mechanism. Sample size $n = 60$ was fixed.

adults.[48] Infants were randomized to ALVAC (n=48) and saline placebo (n=12) and were immunized within 3 days of birth and at 4, 8, and 12 weeks. Several binding antibody measurements were recorded between 10 weeks and 24 months after birth. Our target of inference was average effect of ALVAC vaccination on binding antibodies (bAb) against gp120. Binding antibodies to DP31, a synthetic peptide not included in the vaccine, were also collected in parallel as a proxy for maternal antibody, which were mostly undetectable 6 months after birth. Infants received standard childhood vaccinations including the combination diphtheria tetanus pertussis (DTP) vaccine at 6, 10 and 14 weeks after birth. Antibodies to DTP vaccination were assessed once 6 months after

birth. Immune responses were summarized and analyzed as background blank subtracted optical density (OD) values. We focused on estimating vaccine effects on gp120 antibodies at the Week 10 and 14 visits, which corresponded to assumed peak responses 2 weeks after Doses 3 and 4 of ALVAC respectively. Boxplots of immune responses to DP31 and GP120 over time and scatterplots describing correlations between predictors of the antibody response to gp120 are shown in Figure 4.3.

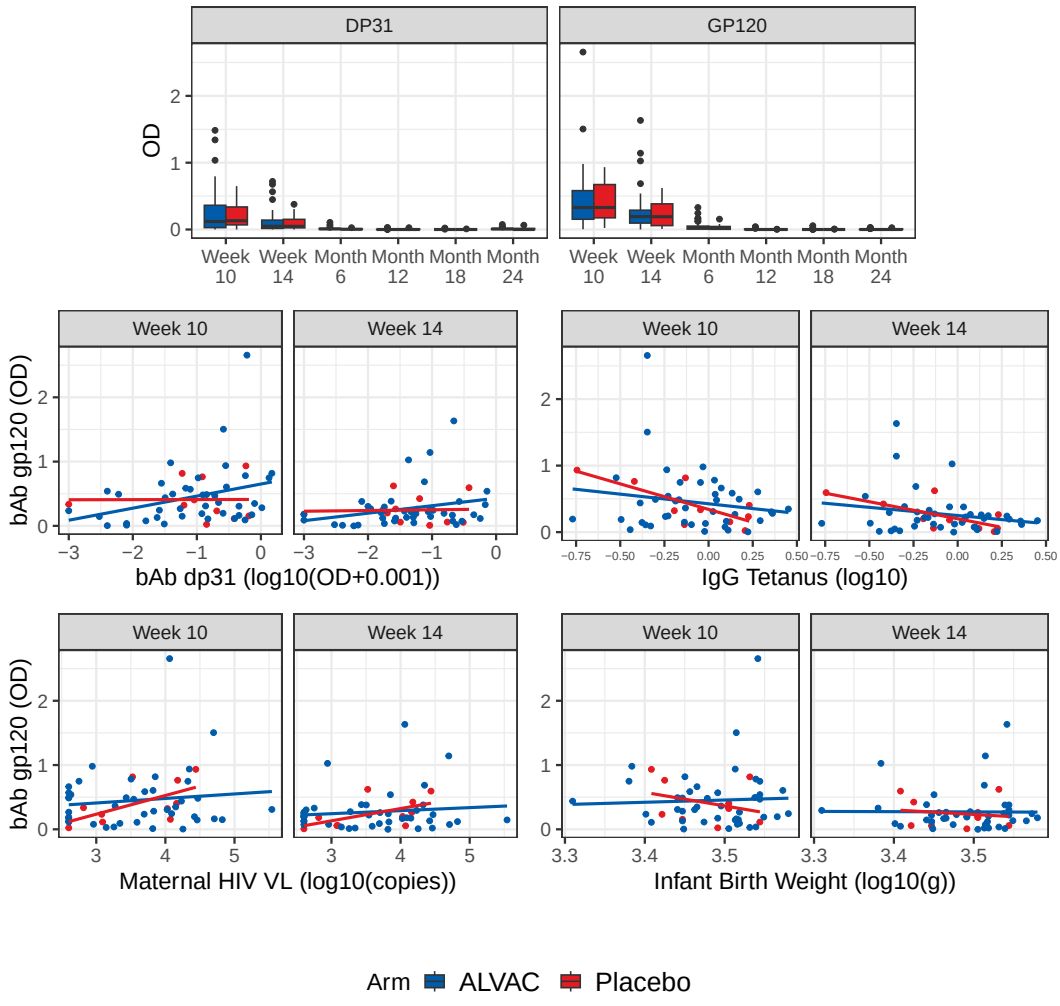

Figure 4.3: Top: boxplots of binding antibodies against synthetic antigen DP31 and target antigen GP120. Below: scatterplots and linear regression models stratified by treatment arm of candidate NCOs (middle rows) and baseline covariates (bottom rows).

We compared the following ATE estimators: (i) plug-in, (ii) covariate-adjusted, (iii) NCO-adjusted, and (iv) a fully-adjusted estimator which adjusted for covariates and NCOs. All estimators used OLS working models fit separately in each treatment arm and HC3 variance corrections. We adjusted for two continuous baseline covariates: infant mass at birth and maternal HIV RNA PCR during the third trimester of pregnancy. We considered adjustment for two NCOs: the level of synthetic DP31 antibody at the same study visit and the antibody response to tetanus vaccination measured six months after birth. Pretests were not performed, since we believed HIV-1 vaccination would have zero effect on immune responses to synthetic or tetanus antigens. We applied log10 transforms to infant birth weight, maternal HIV RNA, and tetanus vaccination response. We applied a log10 transform to DP31 after adding 0.001 to all entries to avoid issues with zero DP31 responses.

Antigen	Contrast	Estimator	Point Estimate (95% CI) (OD)	Variance	Relative Efficiency	Wald p-val
bAb gp120 (Week 10)	ALVAC v Placebo	Plug-in	0.042 (-0.201,0.285)	0.0154	1.000	0.367
		Covariate-adjusted	0.046 (-0.199,0.291)	0.0156	1.010	0.359
		NCO-adjusted (DP31)	0.045 (-0.194,0.284)	0.0149	0.968	0.362
		NCO-adjusted (Tetanus)	**0.043 (-0.156,0.242)**	**0.0103**	**0.673**	**0.334**
		NCO-adjusted (DP31+Tetanus)	0.044 (-0.167,0.254)	0.0115	0.749	0.344
		Covariate+NCO-adjusted	0.043 (-0.22,0.306)	0.0180	1.170	0.377
bAb gp120 (Week 14)	ALVAC v Placebo	Plug-in	0.025 (-0.146,0.196)	0.0076	1.000	0.389
		Covariate-adjusted	0.022 (-0.151,0.195)	0.0078	1.020	0.385
		NCO-adjusted (DP31)	0.026 (-0.141,0.193)	0.0073	0.952	0.383
		NCO-adjusted (Tetanus)	**0.025 (-0.116,0.166)**	**0.0052**	**0.679**	**0.364**
		NCO-adjusted (DP31+Tetanus)	0.025 (-0.119,0.169)	0.0054	0.706	0.392
		Covariate+NCO-adjusted	0.02 (-0.151,0.191)	0.0076	0.998	0.388

Table 4.1: Estimated effect of HIV-1 ALVAC vaccine on IgG binding antibodies to vaccine-matched gp120 antigen in HPTN 027. Vaccine effects are estimated at Weeks 10 and 14, 2 weeks after vaccine doses 3 and 4. At each time point, six different estimates are provided corresponding to different predictor adjustment strategies. Vaccine effect is defined as the population ATE on the additive OD scale.

Results for the data analysis are shown in Table 4.1. All estimators indicated that ALVAC vaccination was associated with an insignificantly higher average level of binding antibody to gp120 at Weeks 10 and 12. All point estimates were very similar, suggesting minimal bias due to adjustment for invalid or skewed NCOs. However, estimators that adjusted for immune response to tetanus vaccination achieved variance reductions exceeding 30% relative to the plug-in estimator. Adjustment

for antibodies to DP31 yielded more modest precision gains of 4 to 5%. Adjustment for baseline covariates led to a loss in precision relative to the naive estimator. While large-sample theory supported the optimality of fully adjusted estimators, the estimator which adjusted for all predictors yielded lower precision, highlighting the penalty paid for adjusting for too many predictors in small samples.

4.4.2 PACTG 230

The Pediatric AIDS Clinical Trials Group (PACTG) 230 trial[8,49] enrolled HEU infants in the United States to determine the safety and immunogenicity of two recombinant gp120 subunit protein vaccines: gp120-MN adsorbed into alum adjuvant (VaxGen) and rgp120-SF2 with MF59 adjuvant (Chiron). Infants were randomized to four injections of VaxGen rgp120 with alum (n=49), Chiron rgp120 with MF59 (n=48), and either placebo or adjuvant alone (n=19) between 0 and 20 weeks of age. For the purposes of our analysis, we grouped together all dose groups and dosing regimens within each platform. HIV Env-specific IgG were measured against HIV Env antigens at week 0, week 10, week 24, week 52, week 76, and week 104 after birth by BAMA. All outcomes were measured as blank-subtracted Mean Fluroescence Intensity (MFI) values. We focused on the week 24 immunogenicity visit corresponding to 4 weeks after the final dose of vaccine and focused on binding antibody IgG against the MN-gp120 antigen. The outcome was log10(MFI+1) transformed. We estimated the population ATE of vaccine on the primary outcome using the following estimators of the vaccine effect: (i) plug-in, (ii) covariate-adjusted, (iii) NCO-adjusted, and (iv) "fully-adjusted" estimator which adjusted for covariates and NCOs. We adjusted for a single baseline covariate: the baseline bAb level to gp120. We adjusted for a single negative control outcome: the contemporaneous level of IgG to gp41, an HIV-1 specific antigen not targeted by the Chiron or VaxGen vaccines. We employed an empirical quantile transform to baseline MN-gp120 IgG and to Week 24 gp41 IgG. All estimators used OLS working models fit separately in each treatment arm and HC3 variance corrections.

Boxplots of the binding antibody responses to the primary antigen (MN-gp120) and negative control antigen (gp41) are shown in the left panel of Figure 4.4. At baseline, most of the outcomes were at an upper limit/saturation point of the assay. Chiron and VaxGen vaccines generated higher

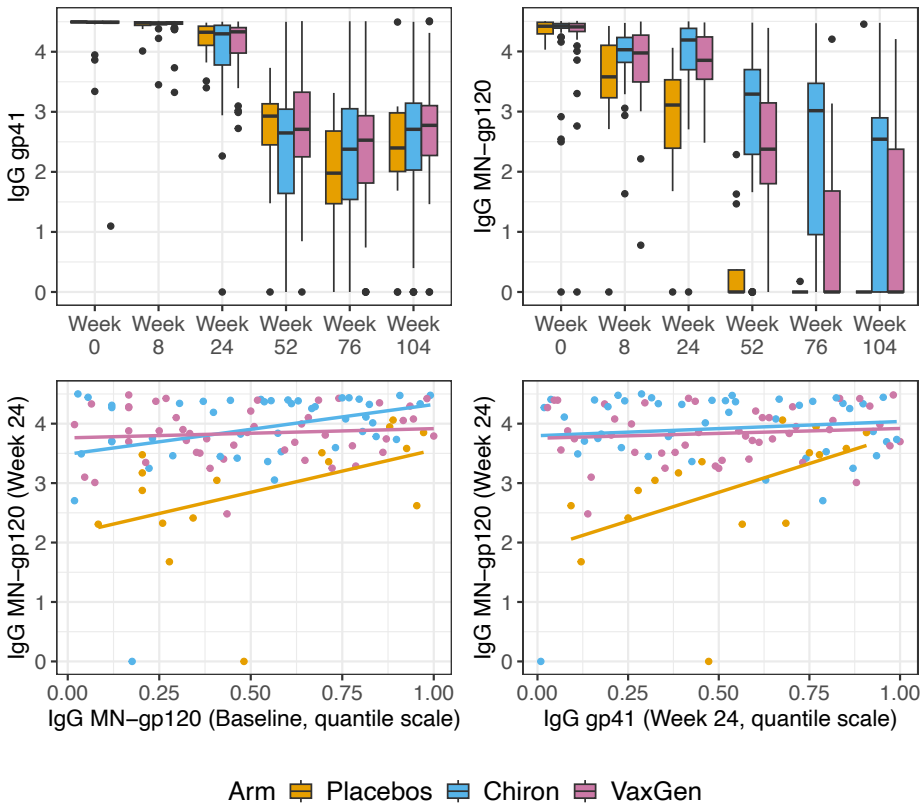

Figure 4.4: Above: boxplots of bAb against vaccine-untargeted HIV-1 antigen gp41 and vaccine targeted HIV-1 antigen MN gp120. Below: association between predictors and bAb to gp120 four weeks after Dose 4 of vaccine/placebo. OLS regressions within each arm shown.

bAb responses to gp120 relative to placebo. The gp41 response appeared very similar between the dose groups, supporting its use as a NCO. Furthermore, when outcomes were plotted against predictors, we observed that bAb to gp120 was positively correlated with the quantile of baseline bAb and quantile of Week 24 bAb. The positive association was most evident in the placebo arm.

Despite the fact that gp41 was not included in either vaccine construct, we wanted to empirically test the NCO assumption prior to adjusting for the NCO. We performed a pretest of the NCO assumption $H_0^{\text{sharp}} : N(1) = N(0)$ on the bAb gp41 responses four weeks after the third dose using adjusted randomization tests. We adjusted for the baseline bAb response to gp120 and used the robust-t statistic based on Lin's adjusted estimator as the test statistic. We simulated the null distribution using 1000 random permutations of the treatment assignment vector. The pretest failed to reject the NCO assumption, as all randomization inference p-values were > 0.85.

Antigen	Contrast	Estimator	Point Estimate (95% CI) (OD)	Variance	Relative Efficiency	Wald p-val
bAb GP120 (4 weeks post Dose 3)	Chiron v Placebo	Plug-in	1.03 (0.48,1.57)	0.0779	1.000	1.16e-04
		Covariate-adjusted	1.02 (0.51,1.53)	0.0688	0.883	5.12e-05
		NCO-adjusted	1.06 (0.55,1.57)	0.0676	0.868	2.19e-05
		Covariate+NCO-adjusted	**1.05 (0.55,1.55)**	**0.0657**	**0.843**	**2.09e-05**
	VaxGen v Placebo	Plug-in	0.95 (0.43,1.47)	0.0698	1.000	1.66e-04
		Covariate-adjusted	0.97 (0.49,1.46)	0.0617	0.884	4.54e-05
		NCO-adjusted	**0.98 (0.51,1.45)**	**0.0575**	**0.824**	**2.35e-05**
		Covariate+NCO-adjusted	0.99 (0.52,1.46)	0.0579	0.830	2.04e-05
	Chiron v VaxGen	Plug-in	0.08 (-0.17,0.33)	0.0166	1.000	2.68e-01
		Covariate-adjusted	0.06 (-0.19,0.31)	0.0163	0.976	3.14e-01
		NCO-adjusted	0.08 (-0.18,0.34)	0.0172	1.030	2.70e-01
		Covariate+NCO-adjusted	0.06 (-0.19,0.32)	0.0168	1.010	3.15e-01

Table 4.2: Estimated effect of gp120 vaccines on IgG binding antibodies to MN-gp120 antigen in PACTG230. Vaccine effects are estimated at Week 24, four weeks after final dose of study vaccines. Three contrasts are given, comparing Chiron and VaxGen vaccines to placebo and the two experimental vaccines head-to-head. For each contrast, four different estimates are provided corresponding to different adjustment strategies. Vaccine effect is defined as the population ATE on the log10(MFI+1) scale.

Results of the data analysis are shown in Table 4.2. When comparing the Chiron and VaxGen vaccines to placebo, adjustment for baseline gp120 bAb led to efficiency improvements of approximately 12% relative to the plug-in. Adjusting for Week 24 bAb to gp41 led to meaningful improvements in efficiency of approximately 13 to 18%. Adjusting for both predictors led to improvements in precision between 16 to 17%. All estimators had very similar point estimates, alleviating concern

of bias due to skewed predictor distributions. However, adjustment for the predictors resulted in no clear improvements in efficiency when comparing the vaccines head-to-head.

4.5 Discussion

We highlight that adjustment for post-baseline negative control outcomes is an useful strategy to sharpen inference on treatment effects in early-phase randomized trials. NCO adjustment is particularly suitable for early-phase trials and secondary analyses, where resources are limited and a primary goal is to improve precision and power. From an analyst perspective, negative control outcomes may offer practical benefits when planning a strategy for adjustment, due to higher correlations with the primary outcome, simplification of the variable selection problem, and ability to capture hard-to-measure characteristics, unfolding developmental processes, or technical sources of variation that cannot be explained by baseline covariates. While large-sample theory supports the optimality of "fully-adjusted" estimators which adjust for all available covariates and NCOs, simulations and application to two early-phase HIV-1 vaccine studies demonstrate that parsimonious adjustment for a limited set of NCOs can improve upon the performance of unadjusted, baseline covariate-adjusted estimators, and fully adjusted estimators in finite samples.

The main drawback of NCO adjustment is the possibility of introducing post-treatment selection bias. NCO-adjustment should only be considered when a candidate NCO is highly unlikely to be affected by the intervention based on subject-matter knowledge or prior experimentation. However, such information may be limited in early-phase trials. We argue that the NCO assumption can be justified in some cases, particularly in vaccine trials, due to the remarkable specificity of the adaptive immune response and complete analyst knowledge of the vaccine construct. If an analyst wishes to check their assumption, we propose pretests of either a sharp null hypothesis or an equivalence hypothesis to guard against violations of the assumption. The sharp null pretest can be used to identify NCO assumption violations in trials that are reasonably well-powered. If the trial is underpowered, an equivalence pretest with an appropriately chosen threshold may be preferred. In Appendix C.3, we offer a simple sensitivity analysis based on linear structural mean models for the potential outcomes $Y(a)$ (which assume that the average treatment effect of A on Y is homogeneous in N). Sensitivity of experimental results to violations of the NCO assumption can be obtained by varying the unknown sensitivity parameter, the average treatment effect of A on N, over a plausible

grid of values. In future vaccine studies, coordinated data collection on a panel of candidate NCOs across multiple studies could help screen for NCOs that are unaffected by vaccination but are prognostic for immune responses of interest.

We primarily focus on treatment effect estimators which rely on linear working models with full predictor by treatment interactions fit using ordinary least squares, owing to their favorable theoretical properties. Recent investigations have proposed flexible, data-adaptive approaches for covariate adjustment.[31,50,51] However, the application of such methods in early-phase trials is yet unclear. Supplying flexible machine learning estimators as working models will require sample-splitting to ensure Type I error rate control,[31,51] which may not be feasible when trial sizes are small and randomization probabilities are imbalanced. The application of flexible models pretrained on historical data or variable selection algorithms such as stepwise selection or LASSO warrant further investigation in small studies. While randomization inference offers an opportunity to leverage data-adaptive methods for efficiency gain, it faces challenges, such as the questionable scientific importance of the sharp null. Randomization tests also do not immediately lend themselves to interval estimators of an interpretable treatment effect parameter. Recent work has used randomization test inversion to develop confidence intervals for quantiles of individual treatment effects,[38] which may be particularly relevant to vaccine studies.[52] Whether these approaches can be extended to incorporate adjustment for covariates and NCOs, perhaps even using flexible data-adaptive methods, is a possible direction for future work.

In principle, we emphasize that negative controls can serve as proxies of the treated or treatment-free potential outcome. In HPTN027, antibodies to the synthetic antigen DP31 are a proxy for the level of maternal antibody and hence, the hypothetical HIV-1-specific antibody response under placebo (i.e., the treatment-free potential outcome). Adjustment for antibodies to DP31 did not reduce variance of the ATE, likely due to the lack of variability in gp120 responses in the trial's placebo arm. In contrast, antibody responses to tetanus vaccination can be considered as a proxy for an infant's hypothetical immune response under assignment to HIV-1 vaccination (i.e., the treated potential outcome).[7] Since infant immune responses in the vaccine arm were more variable, adjustment for tetanus antibodies reduced the variance of the ATE substantially. In PACTG230, antibodies to gp41 were proxies for maternal antibody and the HIV-1-specific antibody response under hypothetical assignment to placebo. Adjustment for the gp41 antibody response reduced the

variance of the ATE substantially when comparing each vaccine to placebo, but not when comparing vaccines head-to-head. Hence, our results suggest that the benefit of NCO adjustment may depend on whether the NCO is a proxy for the treated or treatment-free potential outcome, which causal contrast is of primary interest (e.g., placebo vs. vaccine or comparing several vaccines), and whether the vaccine or placebo arm is expected to contribute more variability to the ATE.

While our work primarily focuses on small, early-phase trials measuring continuous immune response endpoints, improving efficiency of late-phase vaccine efficacy trials via adjustment for off-target infection endpoints warrants further exploration.[53] Our work has implications for study design of future vaccine trials in populations with prior exposure. We advocate for early-phase vaccine immunogenicity studies to measure a variety of auxiliary immune responses believed to be unaffected by the intervention – such as baseline humoral immune responses, concurrent immune responses to vaccine-untargeted antigens, and immune responses to irrelevant vaccinations. Pooling data from placebo arms across several trials to identify valid NCOs and model relationships between auxiliary immune responses and immune responses of interest could hone the efficiency of future trials in the design and analysis stages.

Bibliography

[1] World Health Organization. HIV Statistics, Globally and by WHO Region, 2025. `https://cdn.who.int/media/docs/default-source/hq-hiv-hepatitis-and-stis-library/who-ias-hiv-statistics_2025-new.pdf`, 2025. Information sheet. Accessed May 7, 2026.

[2] Paolo G. Miotti, Taha E. T. Taha, Newton I. Kumwenda, Robert Broadhead, Laban A. R. Mtimavalye, Len Van der Hoeven, John D. Chiphangwi, George Liomba, and Robert J. Biggar. HIV Transmission Through Breastfeeding: A Study in Malawi. *JAMA*, 282(8):744–749, August 1999. ISSN 0098-7484. doi: 10.1001/jama.282.8.744. URL `https://doi.org/10.1001/jama.282.8.744`.

[3] Jeanne Bertolli, Michael E. St. Louis, Robert J. Simonds, Phillip Nieburg, Munkolenkole Kamenga, Christopher Brown, Manzila Tarande, Thomas Quinn, and Chin-Yih Ou. Estimating the timing of mother-to-child transmission of human immunodeficiency virus in a breast-feeding population in kinshasa, zaire. *J Infect Dis*, 174, 1996.

[4] Ria Goswami, Stella J. Berendam, Shuk Hang Li, Ashley N. Nelson, Kristina De Paris, Koen K. A. Van Rompay, Sallie R. Permar, and Genevieve G. Fouda. Harnessing early life immunity to develop a pediatric HIV vaccine that can protect through adolescence. *PLoS Pathog*, 16 (11), 2020. ISSN 1553-7374. doi: 10.1371/journal.ppat.1008983. URL `https://journals.plos.org/plospathogens/article?id=10.1371/journal.ppat.1008983`.

[5] Genevieve G. Fouda, Kristina De Paris, Ofer Levy, Arnaud Marchant, Glenda Gray, Sallie Permar, Mary Marovich, and Anjali Singh. Immunological mechanisms of inducing HIV immunity in infants. *Vaccine*, 38(3):411–415, January 2020. ISSN 1873-2518. doi: 10.1016/j.vaccine.2019.11.011.

[6] Cassandra A. Simonich, Katherine L. Williams, Hans P. Verkerke, James A. Williams, Ruth Nduati, Kelly K. Lee, and Julie Overbaugh. HIV-1 Neutralizing Antibodies with Limited Hypermutation from an Infant. *Cell*, 166(1):77–87, June 2016. ISSN 0092-8674, 1097-4172. doi: 10.1016/j.cell.2016.05.055. URL `https://www.cell.com/cell/abstract/S0092-8674(16)30658-4`. Publisher: Elsevier.

[7] Dean Follmann. Augmented designs to assess immune response in vaccine trials. *Biometrics*, 2006.

[8] GG Fouda, CK Cunningham, EJ McFarland, W Borkowsky, P Muresan, J Pollara, LY Song, BE Liebl, K Whitaker, X Shen, NA Vandergrift, RG Overman, NL Yates, MA Moody, C Fry, JH Kim, NL Michael, M Robb, P Pitisuttithum, J Kaewkungwal, S Nitayaphan, S Rerks-Ngarm, HX Liao, BF Haynes, DC Montefiori, G Ferrari, GD Tomaras, and SR Permar. Infant hiv type 1 gp120 vaccination elicits robust and durable anti-v1v2 immunoglobulin g responses and only rare envelope-specific immunoglobulin a responses. *Journal of Infectious Diseases*, 2015.

[9] Anastasios A. Tsiatis, Marie Davidian, Min Zhang, and Xiaomin Lu. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Stat Med*, 27(23):4658–4677, October 2008. ISSN 0277-6715. doi: 10.1002/sim.3113. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2562926/`.

[10] Ting Ye, Jun Shao, Yanyao Yi, and Qingquan Zhao. Toward better practice of covariate adjustment in analyzing randomized clinical trials. *J Am Stat Assoc*, 2022.

[11] Mia S. Tackney, Tim Morris, Ian White, Clemence Leyrat, Karla Diaz-Ordaz, and Elizabeth Williamson. A comparison of covariate adjustment approaches under model misspecification in individually randomized trials. *Trials*, 24, 2023.

[12] Brennan C Kahan, Vipul Jairath, Caroline J Doré, and Tim P Morris. The risks and rewards of covariate adjustment in randomized trials: an assessment of 12 outcomes from 8 studies. *Trials*, 15(1):139, December 2014. ISSN 1745-6215. doi: 10.1186/1745-6215-15-139. URL `https://trialsjournal.biomedcentral.com/articles/10.1186/1745-6215-15-139`.

[13] Kelly Van Lancker, Frank Bretz, and Oliver Dukes. Covariate adjustment in randomized controlled trials: General concepts and practical considerations. *Clin Trials*, 21, 2024. doi: 10.1177/1740774524125156.

[14] US Food and Drug Administration (FDA). Adjusting for covariates in randomized clinical trials for drugs and biological products: guidance for industry, 2023.

[15] Yunda Huang, Lily Zhang, Holly Janes, Nicole Frahm, Abby Isaacs, Jerome H. Kim, David Montefiori, Julie McElrath, Georgia D. Tomaras, and Peter B. Gilbert. Predictors of durable immune responses six months after the last vaccination in preventive hiv vaccine trials. *Vaccine*, 35, 2017.

[16] Yunda Huang, Yuanyuan Zhang, Kelly E. Seaton, Stephen De Rosa, Jack Heptinstall, Lindsay N. Carpp, April Kaur Randhawa, Lyle R. McKinnon, Paul McLaren, Edna Viegas, Glenda E. Gray, Gavin Churchyard, Susan P. Buchbinder, Srilatha Edupuganti, Linda-Gail Bekker, Michael C. Keefer, Mina C. Hosseinipour, Paul A. Goepfert, Kristen W. Cohen, Brian D. Williamson, M. Juliana McElrath, Georgia D. Tomaras, Juilee Thakar, and James J. Kobie. Baseline host determinants of robust human HIV-1 vaccine-induced immune responses: A meta-analysis of 26 vaccine regimens. *eBioMedicine*, 84, October 2022. ISSN 2352-3964. doi: 10.1016/j.ebiom.2022.104271. URL `https://www.thelancet.com/journals/ebiom/article/PIIS2352-3964(22)00453-4/fulltext`. Publisher: Elsevier.

[17] Ying Huang, Shomoita Alam, Erica Andersen-Nissen, Lindsay N. Carpp, One B. Dintwe, Britta S. Flach, Nicole Grunenberg, Fatima Laher, Stephen C. De Rosa, Guido Ferrari, Craig Innes, Linda-Gail Bekker, James G. Kublin, M. Juliana McElrath, Georgia D. Tomaras, Glenda E. Gray, and Peter B. Gilbert. Non-HIV Vaccine-Induced Immune Responses as Potential Baseline Immunogenicity Predictors of ALVAC-HIV and AIDSVAX B/E-Induced Immune Responses. *Viruses*, 16(9):1365, August 2024. ISSN 1999-4915. doi: 10.3390/v16091365. URL https://www.mdpi.com/1999-4915/16/9/1365.

[18] Peter B Gilbert, Youyi Fong, Michal Juraska, Lindsay N Carpp, Arnold S Monto, Emily T Martin, and Joshua G Petrie. HAI and NAI titer correlates of inactivated and live attenuated influenza vaccine efficacy. *BMC Infect Dis*, 19(1):453, December 2019. ISSN 1471-2334. doi: 10.1186/s12879-019-4049-5. URL https://bmcinfectdis.biomedcentral.com/articles/10.1186/s12879-019-4049-5.

[19] Peter B. Gilbert, Erin E. Gabriel, Xiaopeng Miao, Xiaoming Li, Shu-Chih Su, Janie Parrino, and Ivan S. F. Chan. Fold Rise in Antibody Titers by Measured by Glycoprotein-Based Enzyme-Linked Immunosorbent Assay Is an Excellent Correlate of Protection for a Herpes Zoster Vaccine, Demonstrated via the Vaccine Efficacy Curve. *J Infect Dis*, 210(10):1573–1581, November 2014. ISSN 1537-6613, 0022-1899. doi: 10.1093/infdis/jiu279. URL https://academic.oup.com/jid/article/210/10/1573/2192816.

[20] Paul R. Rosenbaum. The consequences of adjustment for a concomitant variable that has been affected by the treatment. *J R Stat Soc Ser B*, 147, 1984.

[21] Marc Lipsitch, Eric Tchetgen Tchetgen, and Ted Cohen. Negative Controls: A Tool for Detecting Confounding and Bias in Observational Studies. *Epidemiology*, 21(3):383–388, May 2010. ISSN 1044-3983. doi: 10.1097/EDE.0b013e3181d61eeb. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3053408/.

[22] Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. A Selective Review of Negative Control Methods in Epidemiology. *Curr Epidemiol Rep*, 7(4):190–202, December 2020. ISSN 2196-2995. doi: 10.1007/s40471-020-00243-4. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8118596/.

[23] Min Zhang, Anastasios A. Tsiatis, and Marie Davidian. Improving Efficiency of Inferences in Randomized Clinical Trials Using Auxiliary Covariates. *Biometrics*, 64(3):707–715, 2008.

[24] Donald B Rubin. Comment on: Randomization analysis of experimental data in the fisher randomization test, by d basu. *J Am Stat Assoc*, 1980.

[25] M. Elizabeth Halloran, Ira M. Longhini, and Claudio J. Struchiner. Design and interpretation of vaccine field studies. *Epidemiological Reviews*, 21, 1999.

[26] Paul M Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70, 1983.

[27] Donald B Rubin. Bayesian inference for causal effects: The role of randomization. *Ann Stat*, 6, 1978.

[28] Hannah L. Itell, Erin P. McGuire, Petronella Muresan, Coleen K. Cunningham, Elizabeth J. McFarland, William Borkowsky, Sallie R. Permar, and Genevieve G. Fouda. Development and application of a multiplex assay for the simultaneous measurement of antibody responses elicited by common childhood vaccines. *Vaccine*, 36(37):5600–5608, September 2018. ISSN 0264410X. doi: 10.1016/j.vaccine.2018.07.048. URL `https://linkinghub.elsevier.com/retrieve/pii/S0264410X1831017X`.

[29] Walid Kandeil, Miloje Savic, Maria Angeles Ceregido, Adrienne Guignard, Anastasia Kuznetsova, and Piyali Mukherjee. Immune interference (blunting) in the context of maternal immunization with Tdap-containing vaccines: is it a class effect? *Expert Rev Vaccines*, 19(4):341–352, April 2020. ISSN 1476-0584, 1744-8395. doi: 10.1080/14760584.2020.1749597. URL `https://www.tandfonline.com/doi/full/10.1080/14760584.2020.1749597`.

[30] PL Cohen and CB Fogarty. No-harm calibration for generalized Oaxaca–Blinder estimators. *Biometrika*, 111(1):331–338, February 2024.

[31] Marlena S Bannick, Jun Shao, Jingyi Liu, Yu Du, Yanyao Yi, and Ting Ye. A general form of covariate adjustment in randomized clinical trials. *arXiv*, 2023.

[32] James G MacKinnon and Halbert White. Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *J Econometrics*, 29:305–325, 1985.

[33] Joseph E Tota, Frank Struyf, Joshua N Sampson, Paula Gonzalez, Martin Ryser, Rolando Herrero, John Schussler, Naveen Karkada, Ana Cecilia Rodriguez, Nicolas Folschweiller, Carolina Porras, Mark Schiffman, John T Schiller, Wim Quint, Aimée R Kreimer, Cosette M Wheeler, Allan Hildesheim, and Costa Rica Vaccine Trial and PATRICIA Study. Efficacy of the AS04-adjuvanted HPV16/18 vaccine: Pooled analysis of the costa rica vaccine and PATRICIA randomized controlled trials. *J. Natl. Cancer Inst.*, 112(8):818–828, August 2020.

[34] Ethan Ashby, Holly Janes, Dean Follmann, Peter B Gilbert, Honghong Zhou, Xiaowei Wang, Bethany Girard, Frances Priddy, James G Kublin, Lawrence Corey, Kathleen M Neuzil, Lindsey R Baden, Hana M El Sahly, Bo Zhang, and COVE study group. Validating and leveraging non-SARS-CoV-2 respiratory infection as a negative control outcome in a phase 3 COVID-19 vaccine trial with extended observational follow-up. *Am. J. Epidemiol.*, 195(1):168–177, January 2026.

[35] Daniel C Johnson, Elizabeth J McFarland, Petronella Muresan, Terence Fenton, James McNamara, Jennifer S Read, Elizabeth Hawkins, Pamela L Bouquin, Scharla G Estep, Georgia D Tomaras, Carol A Vincent, Mobeen Rathore, Ann J Melvin, Sanjay Gurunathan, and John Lambert. Safety and immunogenicity of an HIV-1 recombinant canarypox vaccine in newborns and infants of HIV-1-infected women. *J. Infect. Dis.*, 192(12):2129–2133, December 2005.

[36] Elizabeth J McFarland, Daniel C Johnson, Petronella Muresan, Terence Fenton, Georgia D Tomaras, James McNamara, Jennifer S Read, Steven D Douglas, Jaime Deville, Marc Gurwith, Sanjay Gurunathan, and John S Lambert. HIV-1 vaccine induced immune responses in newborns of HIV-1 infected mothers. *AIDS*, 20(11):1481–1489, July 2006.

[37] Camila Covián, Ayleen Fernández-Fierro, Angello Retamal-Díaz, Fabián E Díaz, Abel E Vasquez, Margarita K Lay, Claudia A Riedel, Pablo A González, Susan M Bueno, and Alexis M Kalergis. BCG-induced cross-protection and development of trained immunity: Implication for vaccine design. *Front. Immunol.*, 10:2806, November 2019.

[38] Devin Caughey, Allan Dafoe, Xinran Li, and Luke Miratrix. Randomisation inference beyond the sharp null: bounded null hypotheses and quantiles of individual treatment effects. *J R Stat Soc Ser B Stat. Methodol.*, 85(5):1471–1491, February 2024.

[39] Jerzy Neyman. Sur les applications de la theorie des probabilites aux experiences agricoles: Essai des principes. *Stat Sci*, 5, 1921.

[40] Winston Lin. Agnostic notes on regression adjustments to experimental data: Reexamining freedman's critique. *Ann Appl Stat*, 7:295–318, 2013.

[41] Paul R Rosenbaum. Covariance adjustment in randomized experiments and observational studies. *Stat Sci*, 2002.

[42] JW Tukey. Tightening the clinical trial. *Control. Clin. Trials*, 14(4):266–285, August 1993.

[43] MH Gail, WY Tan, and S Piantadosi. Tests for no treatment effect in randomized clinical trials. *Biometrika*, 75(1):57, March 1988.

[44] Jonathan Raz. Testing for no effect when estimating a smooth function by nonparametric regression: A randomization approach. *J. Am. Stat. Assoc.*, 85(409):132, March 1990.

[45] Alisa J Stephens, Eric J Tchetgen Tchetgen, and Victor De Gruttola. Flexible covariate-adjusted exact tests of randomized treatment effects with application to a trial of HIV education. *Ann. Appl. Stat.*, 7(4):2106–2137, December 2013.

[46] Anqi Zhao and Peng Ding. Covariate-adjusted fisher randomization tests for the average treatment effect. *J Econometrics*, 225, 2021.

[47] Pontiano Kaleebu, Harr Freeya Njai, Lei Wang, Norman Jones, Isaac Ssewanyana, Paul Richardson, Kenneth Kintu, Lynda Emel, Philippa Musoke, Mary Glenn Fowler, San-San Ou, Laura Guay, Philip Andrew, Lynn Baglyos, and Huyen Cao team. Immunogenicity of ALVAC-HIV vCP1521 in Infants of HIV-1 Infected Women in Uganda (HPTN 027): the first pediatric HIV vaccine trial in Africa. *Journal of acquired immune deficiency syndromes (1999)*, 65(3):268–277, March 2014. ISSN 1525-4135. doi: 10.1097/01.qai.0000435600.65845.31. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4171956/`.

[48] Supachai Rerks-Ngarm, Punnee Pitisuttithum, Sorachai Nitayaphan, Jaranit Kaewkungwal, Joseph Chiu, Robert Paris, Nakorn Premsri, Chawetsan Namwat, Mark de Souza, Elizabeth Adams, Michael Benenson, Sanjay Gurunathan, Jim Tartaglia, John G. McNeil, Donald P. Francis, Donald Stablein, Deborah L. Birx, Supamit Chunsuttiwat, Chirasak Khamboonruang, Prasert Thongcharoen, Merlin L. Robb, Nelson L. Michael, Prayura Kunasol, and Jerome H. Kim. Vaccination with ALVAC and AIDSVAX to Prevent HIV-1 Infection in Thailand. *N Engl J Med*, 361(23):2209–2220, December 2009. doi: 10.1056/NEJMoa0908492.

[49] Elizabeth J. McFarland, William Borkowsky, Terry Fenton, Diane Wara, James McNamara, Pearl Samson, Minhee Kang, Lynne Mofenson, Coleen Cunningham, Anne-Marie Duliege, Faruk Sinangil, Stephen A. Spector, Eleanor Jimenez, Yvonne Bryson, Sandra Burchett, Lisa M. Frenkel, Ram Yogev, Francis Gigliotti, Katherine Luzuriaga, and AIDS Clinical Trials Group 230 Collaborators. Human immunodeficiency virus type 1 (hiv-1) gp120–specific antibodies in neonates receiving an hiv-1 recombinant gp120 vaccine. *J Infect Dis*, 184, 2001.

[50] Nicholas Williams, Michael Rosenblum, and Iván Díaz. Optimising precision and power by machine learning in randomised trials with ordinal and time-to-event outcomes with an application to COVID-19. *J R Stat Soc Ser A*, page 10.1111/rssa.12915, September 2022. ISSN 0964-1998. doi: 10.1111/rssa.12915. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9539267/`.

[51] Kelly Van Lancker, Iván Díaz, and Stijn Vansteelandt. Automated, efficient and model-free inference for randomized clinical trials via data-driven covariate adjustment, 2024. URL `https://arxiv.org/abs/2404.11150`. Version Number: 1.

[52] Zhe Chen, Xinran Li, and Bo Zhang. The role of randomization inference in unraveling individual treatment effects in clinical trials: Application to HIV vaccine trials. *Stat Commun Infect Dis*, 16(1), 2024.

[53] Lola Etievant, Joshua N. Sampson, and Mitchell H. Gail. Increasing efficiency and reducing bias when assessing HPV vaccination efficacy by using nontargeted HPV strains. *Biometrics*, pages 1–12, 2022. doi: 10.1111/biom.13663.

Chapter 5

IMPROVING THE EFFICIENCY OF INFECTIOUS DISEASE PREVENTION TRIALS USING NEGATIVE CONTROL OUTCOME EVENT TIMES

Summary: *While baseline covariate adjustment can improve the precision of treatment effect estimates in randomized trials, the magnitude of the efficiency gain depends on how prognostic the covariates are for the primary outcome. In randomized trials of vaccines or other immunogens, key prognostic factors for infection risk, such as exposure, can often be conceptualized but are difficult or impossible to measure directly and incorporate into statistical models. We propose adjustment using negative control infections (NCIs), which are assumed to be unaffected by the intervention yet correlated with the primary infection outcome through shared unmeasured prognostic factors. We formalize assumptions under which adjustment for NCIs is valid and examine how right-censoring complicates such adjustment. We introduce an estimator of the treatment-arm–specific survivor functions using the efficient influence function under a model with outcome and covariate subject to right-censoring. We compare the performance of our proposed estimators to standard unadjusted and baseline-covariate-adjusted estimators in numerical experiments. We apply our methods to the HVTN 704 study evaluating the efficacy of a broadly neutralizing antibody for preventing HIV-1 infection, adjusting for time-to-rectal-gonorrhea, an auxiliary outcome unaffected by the experimental intervention and a plausible marker of HIV-1 exposure. While adjustment for baseline covariates only marginally improved precision, adjusting for time-to-rectal gonorrhea improved precision of the log relative risk by 22% relative to the naïve unadjusted analysis. These results highlight the practical benefit of negative control event time adjustment in randomized trials.*

5.1 Introduction

It is well understood that adjusting for baseline covariates can improve efficiency and power in randomized trials by reducing the uncertainty in treatment effect estimates[1–5]. However, the realized benefit depends on how prognostic the set of adjustment variables are for the outcome of interest[6]. If covariates are highly prognostic, substantial efficiency gain can be obtained; if only

weakly prognostic factors are measured, realized precision gains may be negligible.

We focus on the setting of randomized, placebo-controlled trials designed to assess the efficacy of vaccines or other immunogens (e.g., passively administered antibodies) for preventing infectious diseases. In such studies, the prognostic factors determining infection risk can often be conceptualized and qualitatively described but cannot be directly measured and integrated into statistical models[7,8]. For example, many infection outcomes depends on an individual's level of *exposure* to the pathogen of interest. However, quantifying exposure is difficult and in some cases impossible due to practical and conceptual constraints, as we illustrate using two examples below.

1. *Human Immunodeficiency Virus 1 (HIV-1)*: sexual exposure to HIV-1 is quantified by an individual's rate of sexual contacts with virologically unsuppressed partners. While quantifiable in principle, accurately measuring HIV-1 exposure it is very difficult due to reporting biases. Due to stigma, privacy, or personal beliefs, collecting objective data on frequency of sexual contacts, prevention method adherence, and virological suppression statuses of partners is very challenging[9]. Moreover, sexual behaviors which influence HIV-1 exposure may change over time, which may render high-quality baseline measures of exposure only weakly predictive of exposure during study follow-up.

2. *COVID-19*: an individual's SARS-CoV-2 exposure depends on the rate of "contacts" with infected individuals in the community per unit time. We refer to contacts in quotations because for community transmitted respiratory diseases where the only requirement for transmission is spatiotemporal proximity, there is no sufficient criteria to define "contact" between susceptible and infected individuals. Moreover, contact rates will vary over time for a variety of reasons including seasonality, school/work attendance, human mobility, and individual adherence to non-pharmaceutical interventions. Hence, quantifying SARS-CoV-2 exposure is essentially impossible.

Since exposure is frequently difficult or impossible to measure, analysts typically adjust for weakly prognostic but readily accessible baseline predictors (e.g., age and sex), thereby limiting the potential for more efficient randomized vaccine trials. Are there other strategies that may allow us, at least in part, to adjust for the unmeasured factors that we believe determines infection risk?

One idea is to exploit the fact that different infections may be correlated by virtue of their dependence on similar unmeasured factors[10,11]. For example, HIV-1 and other sexually transmitted infections (STIs) depend on an overlapping set of sexual and behavioral characteristics. Hence, other STIs may be compelling markers of exposure to HIV-1 as quantified by the rate of unprotected sexual contacts. This idea is supported by previous studies reporting positive ecological correlations between rates of HIV-1 and other STIs across contexts[12] including in sub-Saharan Africa (SSA)[13] and among populations of men who have sex with men (MSM)[14–17]. Moreover, strong evidence exists linking STI infection to a transient increase in likelihood of HIV-1 acquisition due to mucosal inflammation and increasing viral concentration in the genital tract.[18,19] As another example, it is plausible that SARS-CoV-2 and other community-transmitted respiratory pathogens are correlated through a common set of individual-level characteristics such as mobility, social network connectivity, adherence to non-pharmaceutical interventions, workplace exposure, and others. Positive correlation between SARS-CoV-2 and other pathogens is supported by the abundance of co-infection events[20] and an prior analysis of a pivotal Phase III COVID-19 vaccine efficacy trial showing that incident rhinovirus infection was associated with 3-fold higher odds of experiencing COVID-19 during the same study period.[21]

However, adjusting for auxiliary infection outcomes poses challenges. First, auxiliary infection outcomes are measured post-baseline, and adjustment for post-baseline variables is generally discouraged due to the potential of inducing selection bias in treatment effect estimates. To avoid bias, one must restrict adjustment to auxiliary infection outcomes completely unaffected by vaccination[22,23]. We refer to such vaccine-irrelevant infections as *negative controls infections (NCIs)*[24,25]. Etievant et al. (2023) previously examined if adjusting for observed indicators of NCIs could improve the precision of vaccine efficacy (VE) estimates in Human Papillomavirus (HPV) vaccine trials. However, as we explain in the next section, adjusting for observed NCI indicators relies on a completely random censoring assumption, which is a stronger assumption than needed for many standard unadjusted analyses. Additionally, adjusting for observed NCI indicators represents an extreme coarsening of information that is likely to diminish precision and power gains. Finally, Etievant et al. (2023) did not account for the fact that most infection endpoints in vaccine studies are treated as *time-to-event outcomes*. While there exists many standard methods for analyzing time-to-event outcomes, there are far fewer methods for analyzing data where both the outcome and

a key predictor are both right-censored. We address this gap by proposing an influence-function-based estimator that treats both the outcome and key predictor as right-censored.

In Section 2, we introduce causal assumptions permitting negative control infection adjustment and discuss how adjusting for censored NCIs can lead to assumption violations. In Section 3, we propose an estimator of the treatment-arm-specific survival function of a primary infection outcome using the observed-data efficient influence function (EIF) when the key predictor, the time-to-negative control infection, is also right-censored. In Section 4, we compare our proposed estimators to alternative unadjusted and baseline-covariate-adjusted estimators in numerical experiments under data generating models consistent with randomized vaccine trials. In Section 5, we apply our methods to the HVTN704 trial evaluating the efficacy of VRCO1, a broadly neutralizing antibody, in preventing HIV-1 infection in a MSM population.

5.2 Methods

5.2.1 Notation and Data Structure

Consider a randomized trial which collects fully-observed baseline covariates X (e.g., age) and randomly assigns a binary treatment A (vaccine or placebo). The primary outcome of the study is Y, the time-to-infection caused by target vaccine-preventable pathogen. Through routine surveillance or standard of care, the study also measures N which denotes time-to-infection caused by another auxiliary pathogen. We assume the oracle data units are drawn independent and identically distributed samples from a distribution P_0 in statistical model M.

$$O^{\text{oracle}} = (A, X, Y, N) \overset{iid}{\sim} P_0 \in M$$

In the developments below, we will frequently refer to the "full" data unit, where only the *outcome* is right-censored, but the negative control infection time is fully observed.

$$O^{\text{full}} = (A, X, \tilde{Y}, \Delta_Y, N)$$

Where $\tilde{Y} := \min(Y, C_Y)$ and $\Delta_Y := I(Y \leq C_Y)$, where C_Y is the censoring time for the vaccine-preventable pathogen.

In reality, we observe a data unit where N is also coarsened due to right-censoring. Define the observed data unit as follows.

$$O^{\text{obs}} = (A, X, \tilde{Y}, \Delta_Y, \tilde{N}, \Delta_N)$$

Where $\tilde{N} := \min(N, C_N)$ and $\Delta_N := I(N \leq C_N)$, where C_N is the censoring time for the vaccine-irrelevant pathogen.

In the sequel, we will distinguish between censoring due to dropout (C) and administrative censoring (C_{admin}), as administrative censoring is usually pre-specified in advance of the trial, and can be treated like a baseline covariate. We also do not impose restrictions on the censoring times C_N and C_Y. If Y and N are subject to a common censoring time, $C_N = C_Y$ with probability 1. In other cases, when N and Y are measured at different time grids, C_N and C_Y may differ.

5.2.2 Estimand, Causal Assumptions, and Identification

We adopt Neyman and Rubin's potential outcomes framework[27,28], which presumes that each participant has a tuple of potential outcomes, $\{Y(0), Y(1)\}$. Let $S_a(t) = P_0(Y(a) > t)$ refer to the treatment-arm-specific survivor function with respect to the vaccine-preventable outcome of interest. Our interest lies in identifying $S_a(t_0)$ for $a \in \{0, 1\}$ at some landmark time t_0, such that we can construct contrasts which quantify the vaccine efficacy. For example, we quantify the treatment effect using the *causal log cumulative incidence ratio* at time t_0.

$$\log \text{RR}(t_0) = \log \left(\frac{1 - S_1(t_0)}{1 - S_0(t_0)} \right) \tag{5.1}$$

5.1 quantifies the (log) relative cumulative infection risk at time t_0 caused by vaccination. One minus the exponentiated version of 5.1 corresponds to the *prevention efficacy (PE)* at time t_0. We choose to focus on this cumulative-incidence-based estimand, as it has a clear causal interpretation without relying on a statistical model for validity (such as the hazard ratio).[29] We can interpret $\text{PE}(t_0) = X\%$ as indicating a vaccinated individual would reduce their probability of infection by time t_0 by $X\%$ on average relative to unvaccinated individuals.

We impose causal assumptions that identify the target parameter 5.1 from observed data. The

assumptions can be summarized graphically in the directed acyclic graph (DAG) shown in Figure 5.1.

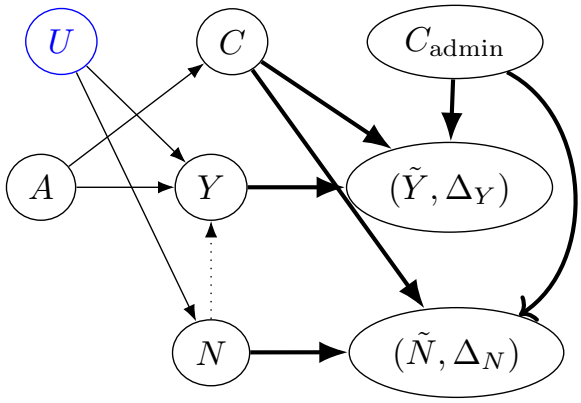

Figure 5.1: Graphical causal model illustrating assumptions underlying approach. U is an unmeasured variable that is a common cause of both Y and N that is presumed to drive correlation between the infection types. The figure assumes a common dropout time $C = C_Y = C_N$. Arrows indicate causal relationships while while **bold** arrows denote deterministic relationships (implied by the deterministic functions $\min(x, y)$ and $I(x \leq y)$ used to obtain the minimum event time and the event indicator)

Assumption 17 (Stable Unit Treatment Value Assumption (SUTVA))**.** *SUTVA consists of the assumptions of (i) consistency of potential outcomes, (ii) no interference between units, (iii) no multiple versions of treatment*[30]

Consistency of potential outcomes and no multiple versions of treatment are standard assumptions that are satisfied in most randomized vaccine trials. In infectious disease settings, the assumption of no interference between units may be violated due to spillover or herd immunity effects. This is a concern when the trial enrolls clusters of participants – such as partners, households, or communities – where the treatment status of one unit can affect the outcome of another unit. However, in individually randomized vaccine trials, participants are embedded in a much larger target population such that contacts between study participants are very unlikely. We consider individually randomized vaccine trials where the impact of interference can be safely assumed to be negligible.

Assumption 18 (Positive treatment probability)**.**

$$0 < \pi_0(a \mid x) \;\; \textit{for } a \in \{0,1\} \textit{ and all } x \textit{ with positive support}$$

Where $\pi_0(a \mid x) := P_0(A = a \mid X = x)$ is the conditional probability of receiving treatment $A = a$. Positivity implies that treatment is nondeterministic, which is satisfied by design in randomized trials.

Assumption 19 (Strong ignorability[30])**.**

$$A \perp (Y(0), Y(1), N(0), N(1), C_N(0), C_N(1), C_Y(0), C_Y(1), X)$$

Where $(N(0), N(1))$ are the potential outcomes for off-target infection, and $(C_Y(1), C_Y(0))$ and $(C_N(0), C_N(1))$ are the vector of potential censoring times for Y and N respectively.

Assumption 17 to 19 are sufficient to identify $S_a(t_0)$ under the *oracle data*, O^{oracle}. However, if patients infection times are right-censored, additional assumptions on the censoring mechanism are required to identify the target parameter from O^{full} and O^{obs}.

Assumption 20 (Independent censoring)**.**

$$(C_Y, C_N) \perp (Y, N) \mid (A, X)$$
$$P_0(C_Y \geq t_0 \mid A, X) \geq \epsilon > 0 \qquad P_0(C_N \geq \min(\tilde{N}, t_0) \mid A, X, \tilde{Y}, \Delta_Y) \geq \epsilon > 0$$

The first part of Assumption 20 assumes that the censoring times are independent of the event times conditional on baseline factors, an assumption frequently made in survival analysis. The latter case ensures that censoring is not exhaustive within subgroups of participants at the relevant evaluation points: at the landmark time t_0 for C_Y and at each participant's observed $\tilde{N} \wedge t_0$ for C_N. In settings where N and Y share a common censoring mechanism (i.e., $C_Y = C_N$), when $\Delta_Y = 0$, $P_0(C_N \geq t \mid a, x, \tilde{y}, \delta_y) = I(t \leq \tilde{y})$, which still satisfies positivity at the relevant evaluation point. Some important consequences of Assumption 20 are established in Remark 1 below. A proof of this result can be found in Appendix D.1.1.

Remark 1. *Two important implications of Assumption 20 are the following*

$$C_Y \perp Y \mid (A, X) \qquad C_N \perp N \mid (A, X, \tilde{Y}, \Delta_Y)$$

While not required for identification, we introduce other assumptions under which adjusting for the auxiliary infection time N avoids post-treatment selection bias while unlocking the potential for more efficiency gains.

Assumption 21 (Negative control infection (NCI))**.**

$$N = N(0) = N(1) \quad \textit{with probability } 1$$

The negative control assumption requires that N satisfy the sharp causal null hypothesis of no treatment effect. By precluding an effect of A on N, we can adjust for N without incurring selection bias[23]. A consequence of the NCI assumption is statistical independence between $N \perp A$ conditional on any other variable(s). Next, we introduce an assumption that is sufficient to describe cases where adjusting for N can lead to precision gains.

Assumption 22 (U-comparability[25])**.** *Let U be an unmeasured cause of $Y(a)$. For example, suppose $Y(a)$ and N are generated from the nonparametric structural equation model (NPSEM)*

$$Y(a) = h_Y(U, A = a, X, N, \epsilon_Y)$$
$$N = h_N(U, X, \epsilon_N)$$

Where h_Y, h_N are unknown deterministic function and $\epsilon_Y \perp \epsilon_N$ are completely random errors. Under the NPSEMs, N is correlated with U conditional on the other measured variables.

$$N \not\perp U \mid A, X$$

Assumption 21 assumes that there exists a unmeasured variable U that influences both Y and N. In infectious disease settings, one can consider U to capture unmeasured or fundamentally unmeasurable factors such as "exposure" which are believed to be similar for both infection outcomes.

As depicted on the DAG in Figure 5.1, the structural equation for N in Assumption 22 does not depend on the treatment variable A. The structural equation for $Y(a)$ *may* depend on N but not vice versa, otherwise N may be influenced by treatment indirectly via its affect on Y.

Following the arguments of Westling et al. (2023), Assumptions 17-22 are sufficient to identify $S_a(t)$ from the *full data*, O^{full}, using the G-formula applied to the conditional treatment-arm-specific survivor function.

$$S_{a_0}(t) = \mathbb{E}_{O^{\text{full}}}[S_{0Y}(t \mid a_0, X, N)]$$

Where $S_{0Y}(t \mid a_0, X, N) := P_0(Y > t \mid A = a_0, X, N)$.

Next we consider identification using the *observed data*, O^{obs}, where N may be right-censored. We introduce the following proposition, which is proven in the Appendix.

Proposition 4. *Let $O^{obs} := (A, X, \tilde{Y}, \Delta_Y, \tilde{N}, \Delta_N)$. Under assumption 17 to 22, the following G-formulas that attempt to adjust for the coarsened versions of N, $\tilde{N}$ and Δ_N respectively, do not identify $S_a(t)$ due to post-treatment selection.*

$$\mathbb{E}_{O^{obs}}\left[P(Y > t \mid A = a_0, X, \tilde{N})\right] \neq S_{a_0}(t)$$
$$\mathbb{E}_{O^{obs}}\left[P(Y > t \mid A = a_0, X, \Delta_N)\right] \neq S_{a_0}(t)$$

Rather, we can identify $S_a(t)$ using inverse probability weighting (IPW) of the treated, complete-N-cases.

$$\mathbb{E}_{O^{obs}}\left[\frac{\Delta_N I(A = a_0)}{G_{0N}(\tilde{N} \mid A, X, \tilde{Y}, \Delta_Y)\pi_0(a_0 \mid X)} S_{0Y}(t \mid a_0, X, N = \tilde{N})\right] = S_{a_0}(t)$$

In short, as described by Rosenbaum (1984), adjusting for observed versions of the NCI – $\tilde{N}$ or Δ_N – will fail to identify $S_a(t)$ if censoring is influence by treatment. Assuming censoring times are completely unaffected by treatment may be plausible if participants are blinded and the treatment is not associated with enhanced/reduced likelihood of study dropout. However, this assumption is hard to justify through scientific knowledge and is a stronger assumption than required for a standard Kaplan-Meier analysis. We deem this assumption unnecessarily strong. Our

identification approach focuses adjustment on the true negative control infection time, N, which is partially observed due to right-censoring.

In the next section, we will use the IPW identification approach in 4 as the basis for constructing an estimating equation for $S_{a_0}(t_0)$ with favorable theoretical properties.

5.2.3 Efficiency calculations

In this section, we will focus on the derivation of the efficient influence function (EIF) of the treatment-arm-specific survival probability $S_{a_0}(t_0)$ under the full and observed data generating models. We focus on the EIF because it can be used to construct the most efficient estimator of $S_{a_0}(t_0)$ among all asymptotically linear estimators. Then, inference on $S_a(t_0)$ for $a \in \{0, 1\}$ can be translated to inference on the target parameter $\log \mathrm{RR}(t_0)$ using the delta method.

Following Theorem 1 of Westling et al. (2023), in the model where O^{full} is observed, the EIF of $S_a(t_0)$ is given by

$$\phi^F_{a_0,t_0}(o^{\text{full}}) = S_{0Y}(t_0 \mid a_0, x, n)\left[1 - \frac{I(a = a_0)}{\pi_0(a_0 \mid x)}\left\{\frac{I(\tilde{y} \le t_0, \delta_y = 1)}{S_{0Y}(\tilde{y} \mid a, x, n)G_{0Y}(\tilde{y} \mid a, x)} - \int_0^{\min(t_0,\tilde{y})} \frac{\Lambda_{0Y}(du \mid a, x, n)}{S_{0Y}(u \mid a, x, n)G_{0Y}(u \mid a, x)}\right\}\right] \tag{5.2}$$

Where $S_{0Y}(t \mid a, x, n) := P_0(Y > t \mid A = a, X = x, N = n)$ is the conditional survival function for the primary outcome, $\Lambda_{0Y}(u \mid a, x, n)$ is the corresponding cumulative hazard, $G_{0Y}(t \mid a, x) := P_0(C_Y \ge t \mid A = a, X = x)$ is the conditional survival function of C_Y, and $\pi(a \mid x) := P_0(A = a \mid X = x)$ is the treatment propensity.

In practice, N is subject to right-censoring, and we observe

$$O^{\text{obs}} = (A, X, \tilde{Y}, \Delta_Y, \tilde{N}, \Delta_N) \qquad \tilde{N} := \min(N, C_N) \qquad \Delta_N := I(N \le C_N)$$

We make no assumption about the joint relationship between C_N and C_Y. They may coincide when N and Y are measured on a common visit grid and are subject to the same dropout process, they may be deterministically related, or they differ in arbitrary ways. The only restriction we place is independent censoring (Assumption 20) which as described in Remark 1, implies the following

coarsening at random assumption for N.

$$C_N \perp N \mid (A, X, \tilde{Y}, \Delta_Y) \tag{5.3}$$

Before we present the observed-data efficient influence function, we define two key nuisance parameters. The first is the conditional survival function of C_N.

$$G_{0N}(t \mid a, x, \tilde{y}, \delta_y) := \exp\left\{-\int_0^t \lambda_{C_N}(u \mid a, x, \tilde{y}, \delta_Y)du\right\}$$

We note that when both Y and N are subject to the same censoring time ($C = C_N = C_Y$), G_{0N} is a derived quantity from G_{0Y} rather than a separate nuisance. $G_{0N}(\cdot \mid a, x, \tilde{y}, \delta_y = 0)$ yields a point mass at $\tilde{y}$ while when $G_{0N}(\cdot \mid a, x, \tilde{y}, \delta_y = 1) \equiv G_{0N}(\cdot \mid a, x, C > \tilde{y})$ yields a left-truncated survival function.

The second key nuisance is the conditional mean of the full-data EIF, conditional on all the other observed variables.

$$Q_0(u \mid a, x, \tilde{y}, \delta_y) := \mathbb{E}[\phi^F_{a_0,t_0}(\tilde{y}, \delta_y, a, x, N) \mid N \geq u, A = a, X = x, \tilde{Y} = \tilde{y}, \Delta_Y = \delta_y]$$

By applying the theory of influence function projections in coarsened data models,[32], we obtain the observed-data EIF.

Theorem 1 (Observed-data EIF)**.** *Let $\phi^F_{a_0,t_0}$ denote the full-data EIF as shown in 5.2. Under censoring at random (Assumption 4), the observed-data influence function corresponding to $O^{obs} = (A, X, \tilde{N}, \Delta_N, \tilde{Y}, \Delta_Y)$ is given by the following.*

$$\phi^{obs}_{a_0,t_0}(o^{obs}) = \frac{\delta_n}{G_{0N}(\tilde{n} \mid a, x, \tilde{y}, \delta_y)}\phi^F_{a_0,t_0}(\tilde{y}, \delta_y, a, x, \tilde{n}) + \int_0^\infty \frac{Q_0(u \mid a, x, \tilde{y}, \delta_y)}{G_{0N}(u \mid a, x, \tilde{y}, \delta_y)}dM_{C_N}(u \mid a, x, \tilde{y}, \delta_y) \tag{5.4}$$

Where $dM_{C_N}(u \mid a, x, \tilde{y}, \delta_y) := dN_{C_N}(u) - \lambda_{C_N}(u \mid a, x, \tilde{y}, \delta_y)I(N \geq u, C_N \geq u)du$ is the canonical martingale increment for the censoring variable C_N, where $dN_{C_N}(u) = I(C_N = u, N > u)$ is the jump process for NCO censoring and $\lambda_{C_N}(u \mid a, x, \tilde{y}, \delta_y)$ is the NCO censoring hazard.

Moreover, because $\phi^F_{a_0,t_0}$ is the nonparametric efficient influence function in the full-data model

and $\phi^{obs}_{a_0,t_0}$ is its unique projection onto the observed-data tangent space, $\phi^{obs}_{a_0,t_0}$ is necessarily the observed-data efficient influence function.

Theorem 1 is proved in Appendix D.2. Our interest in the observed-data efficient influence function (EIF) stems from its role as a optimal estimating function for the target parameter $S_a(t)$. Next, we will introduce a one-step estimator for $S_{a_0}(t)$ built from the EIF.

5.2.4 A cross-fitted one-step estimator

As written, $\phi^{\text{obs}}_{a_0,t_0}$ depends on at least four nuisance parameters: the treatment propensity $\pi_0(a \mid x)$, the conditional survival function of Y-censoring $G_{0Y}(t \mid a, x)$, the conditional density of the irrelevant infection outcome $f_{0N}(n \mid a, x, \tilde{y}, \delta_y)$, and the conditional survival function of the primary outcome $S_{0Y}(y \mid a, x, n)$. If N and Y are subject to different censoring processes, we need to estimate an additional nuisance parameter, the conditional survival function of C_N: $G_{0N}(t \mid a, x, \tilde{y}, \delta_y)$. Given estimators of the nuisances, there are many potential estimators of the $S_{a_0}(t_0)$. Let $\hat{\phi}^{\text{obs}}_{n,a_0,t_0}$ denote the value of $\phi^{\text{obs}}_{a_0,t_0}$ when the true unknown nuisances are replaced by arbitrary estimators. The classic one-step estimator would be $\mathbb{P}_n\hat{\phi}^{\text{obs}}_{n,a_0,t_0}$, where $\mathbb{P}_n$ denotes the empirical probability measure (i.e., sample average).[32,33] In this case, the one-step estimator is also an estimation equations-based estimator because the influence function is linear in the target parameter.

Asymptotically linear estimators, including the one-step estimator, depend on nuisance estimators in two important ways. First, negligibility of the second-order remainder term requires the nuisance estimators to converge sufficiently quickly to their true values. Second, negligibility of empirical process terms typically requires the nuisance estimators to belong to sufficiently small function classes. In practice, correctly specified parametric models for nuisance functions are rarely available, motivating the use of flexible, data-adaptive estimators. Because the remainder term is second-order in the nuisance estimation error, valid inference can often be obtained even when nuisance estimators converge at slower-than-root-n rates. Pairing cross-fitting with data-adaptive nuisance estimation removes the complexity constraint on nuisance estimators required to control the empirical process terms.[34? ?]

Following Westling et al.[31], for a deterministic integer $2 \leq K \leq \lfloor n/2 \rfloor$, we randomly participant $\{1, \ldots, n\}$ into K disjoint folds $\mathcal{V}_{n,1}, \ldots, \mathcal{V}_{n,K}$. For each $k \in \{1, \ldots, K\}$, we define $\mathcal{T}_{n,k} := \{O_i^{\text{obs}} :$

$i \notin \mathcal{V}_{n,k}\}$ as the training set for fold k. Define $\hat{\phi}^{\text{obs}}_{n,k,a_0,t_0}$ as the EIF where nuisance estimated only using the observations from the training set $\mathcal{T}_{n,k}$ are substituted for their true counterparts. We define the cross-fitted one-step estimator for $S_{a_0}(t_0)$ as follows

$$\hat{S}_{a_0}(t) := \frac{1}{n}\sum_{k=1}^{K}\sum_{i\in\mathcal{V}_{n,k}} \hat{\phi}^{\text{obs}}_{n,k,a_0,t_0}(O_i^{\text{obs}}). \tag{5.5}$$

Below, we introduce two important theoretical results for the cross-fitted, one-step estimator. The first establishes that multiple-robustness of the estimator, meaning that the estimator is consistent when only combinations of a subset of nuisance functions are correctly specified. The second establishes the conditions under which the cross-fitted, one-step estimator is asymptotically linear with variance estimator given by the variance of the observed data influence function (and is therefore efficient).

The theoretical results use the following standard regularity conditions:

(C1) There exists π_∞, $G_{\infty,Y}$, $G_{\infty,N}$, $S_{\infty,Y}$, and $f_{\infty,N}$ such that:

$$\begin{aligned}
&\max_k \mathbb{E}\left\{\frac{1}{\widehat{\pi}_{n,k}(a_0 \mid X)} - \frac{1}{\pi_\infty(a_0 \mid X)}\right\}^2 \xrightarrow{p} 0,\\
&\max_k \mathbb{E}\left\{\sup_{u\in[0,t_0]}\left|\frac{1}{\widehat{G}_{n,k,Y}(u \mid a_0, X)} - \frac{1}{G_{\infty,Y}(u \mid a_0, X)}\right|\right\}^2 \xrightarrow{p} 0,\\
&\max_k \mathbb{E}\left\{\sup_{u\in[0,t_0]}\left|\frac{\widehat{S}_{n,k,Y}(t_0 \mid a_0, X)}{\widehat{S}_{n,k,Y}(u \mid a_0, X)} - \frac{S_{\infty,Y}(t_0 \mid a_0, X)}{S_{\infty,Y}(u \mid a_0, X)}\right|\right\}^2 \xrightarrow{p} 0,\\
&\max_k \mathbb{E}\left\{\sup_{u\in[0,\tau_N]}\left|\frac{1}{\widehat{G}_{n,k,N}(u \mid a_0, X, \widetilde{Y}, \Delta_Y)} - \frac{1}{G_{\infty,N}(u \mid a_0, X, \widetilde{Y}, \Delta_Y)}\right|\right\}^2 \xrightarrow{p} 0,\\
&\max_k \mathbb{E}\left\{\sup_{u\in[0,\tau_N]}\left|\widehat{Q}_{n,k}(u \mid a_0, X, \widetilde{Y}, \Delta_Y) - Q_\infty(u \mid a_0, X, \widetilde{Y}, \Delta_Y)\right|\right\}^2 \xrightarrow{p} 0.
\end{aligned}$$

(C2) There exists a constant $\varepsilon > 0$ such that, with probability tending to one, for almost every relevant covariate value, $\widehat{\pi}_{n,k}(a_0 \mid x) \geq \varepsilon$, $\pi_\infty(a_0 \mid x) \geq \varepsilon$, $\widehat{G}_{n,k,Y}(t_0 \mid a_0, X) \geq \varepsilon$, $G_{\infty,Y}(t_0 \mid a_0, x) \geq \varepsilon$, $\widehat{G}_{n,k,N}(u \mid a_0, x, \widetilde{y}, \delta_y) \geq \varepsilon$, and $G_{\infty,N}(u \mid a_0, x, \widetilde{y}, \delta_y) \geq \varepsilon$.

(C3) For almost every relevant covariate value, the limiting nuisance functions satisfy at least one

of the following correctness regimes:

$$(S_{\infty,Y}, G_{\infty,N}) = (S_{0Y}, G_{0N}),$$

or

$$(S_{\infty,Y}, f_{\infty,N}) = (S_{0Y}, f_{0N}),$$

or

$$(\pi_\infty, G_{\infty,Y}, G_{\infty,N}) = (\pi_0, G_{0Y}, G_{0N}),$$

or

$$(\pi_\infty, G_{\infty,Y}, f_{\infty,N}) = (\pi_0, G_{0Y}, f_{0N}).$$

(C4) We define

$$r_{a_0,t_0,1} \equiv \max_k \mathbb{E}\left|\{\widehat{\pi}_{n,k}(a_0 \mid X) - \pi_0(a_0 \mid X)\}\{\widehat{S}_{n,k,Y}(t_0 \mid a_0, X, N) - S_{0Y}(t_0 \mid a_0, X, N)\}\right|,$$
$$r_{a_0,t_0,2} \equiv \max_k \mathbb{E}\left|\widehat{S}_{m,k,Y}(t_0 \mid a_0, X, N)\int_0^{t_0}\left\{\frac{G_{0Y}(u \mid a_0, X, N)}{\widehat{G}_{n,k,Y}(u \mid a_0, X, N)} - 1\right\}\left(\frac{S_{0Y}}{\widehat{S}_{n,k,Y}} - 1\right)(du \mid a_0, X, N)\right|,$$
$$r_{a_0,t_0,3} \equiv \max_k \mathbb{E}\left|\int_0^N \{\widehat{Q}_{n,k}(u) - Q_k(u)\}\{d\widehat{\Lambda}_{n,k,C_N}(u) - d\Lambda_{0C_N}(u)\}\right|.$$

It holds that $r_{a_0,t_0,1} = o_P(n^{-1/2})$, $r_{a_0,t_0,2} = o_P(n^{-1/2})$, and $r_{a_0,t_0,3} = o_P(n^{-1/2})$.

Theorem 2 (Multiple robustness)**.** *Suppose Assumptions 1-6 and regularity conditions C1-C3 hold. Then the cross-fitted one-step estimator $\hat{S}_{n,a_0}(t_0)$ is consistent for $S_{a_0}(t_0)$ provided at least one of the following nuisance-function sets is consistently estimated:*

$$(S_Y, G_N), \qquad (S_Y, f_N), \qquad (\pi, G_Y, G_N), \qquad (\pi, G_Y, f_N).$$

Thus, consistency does not require all nuisance functions to be correctly specified simultaneously.

Theorem 3 (Asymptotic linearity)**.** *Suppose Assumptions 1-6 and regularity conditions C1-C3 hold. Under an additional condition C4, which claims that the nuisance estimators converge quickly enough to their true values such that remainder terms are $o_P(n^{-1/2})$, then*

$$\sqrt{n}\{\widehat{S}_{a_0}(t_0) - S_{a_0}(t_0)\} = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \left\{ \phi_0^{\text{obs}}(O_i^{\text{obs}}) - S_{a_0}(t_0) \right\} + o_p(1).$$

Consequently,

$$\sqrt{n}\{\widehat{S}_{a_0}(t_0) - S_{a_0}(t_0)\} \rightsquigarrow N(0, \sigma_0^2),$$

where

$$\sigma_0^2 = E_0 \left[\{\phi_0^{\text{obs}}(O^{\text{obs}}) - S_{a_0}(t_0)\}^2 \right].$$

These results are proven in Appendices D.3 and D.4.

5.2.5 Nuisance parameter specification and estimation

Recall that the EIF and cross-fitted one-step estimator depend on several nuisance functions that must be estimated from the observed data. In a randomized trial, the treatment propensity score $\pi_0(a \mid x)$ is known by design and can therefore be estimated consistently using a simple logistic regression model. For the remaining nuisance functions, correctly specified parametric models are generally unavailable. We therefore advocate estimating these quantities using cross-fitted SuperLearner ensembles[31,35], which combine predictions from a library of candidate learners and protect against nuisance misspecification.

For the censoring survival functions ($G_{0Y}(t \mid a, x)$ and $G_{0N}(t \mid a, x, \widetilde{Y}, \Delta_Y)$), we recommend SuperLearner libraries containing parametric, semiparametric, and nonparametric survival models, such as Weibull regression, Cox regression, generalized additive Cox models, and random survival forests. When (C_Y) and (C_N) are deterministically related (e.g., ($C_N = C_Y$)), the estimator of (G_{0N}) can be obtained directly from the estimator of (G_{0Y}).

The remaining nuisance functions,

$$f_{0N}(n \mid a, x, \widetilde{Y}, \Delta_Y) \quad \text{and} \quad S_{0Y}(y \mid a, x, n),$$

describe the conditional distribution of one event time given the other. Estimating these quantities separately can lead to incompatibility, meaning that the resulting conditional distributions do not correspond to any valid joint distribution of $(Y, N) \mid (A, X)$ [36,37]. To avoid this issue, we instead model the joint distribution through the one-way factorization of the joint likelihood

$$f_N(n \mid a, x) \times S_{0Y}(y \mid a, x, n).$$

This parameterization guarantees compatibility by construction and yields the required nuisance function ($f_{0N}(n \mid a, x, \widetilde{Y}, \Delta_Y)$) through an application of Bayes' rule which is shown below and proven in Appendix D.5.

Proposition 5 (Deriving target nuisance under one-way factorization)**.** *We can express the target nuisance $f_{0N}(n \mid a, x, \tilde{y}, \delta_y)$ in terms of the one-way factorization $f_N(n \mid a, x)$ and $S_{0Y}(y \mid a, x, n)$ via Bayes rule.*

$$f_{0N}(n \mid a, x, \tilde{y}, \delta_y) = f_N(n \mid a, x) \cdot \frac{L(\tilde{y}, \delta_y \mid a, x, n)}{L(\tilde{y}, \delta_y \mid a, x)}$$

where

$$\frac{L(\tilde{y}, \delta_y \mid a, x, n)}{L(\tilde{y}, \delta_y \mid a, x)} = \left[\delta_y \left(\frac{f_{0Y}(\tilde{y} \mid a, x, n)}{f_Y(\tilde{y} \mid a, x)}\right) + (1 - \delta_y)\frac{S_{0Y}(\tilde{y} \mid a, x, n)}{S_Y(\tilde{y} \mid a, x)}\right].$$

Note that $S_Y(u \mid a, x) = \int_0^\infty S_{0Y}(u \mid a, x, n) f_N(n \mid a, x) dn$ by tower law. f_{0Y} and f_Y refer to the conditional density functions associated with S_{0Y} and S_Y, which can be obtained from the survival functions using the following identity $f(x) = -(d/du)S(u) \mid_{u=x}$.

Hence, $f_{0N}(n \mid a, x, \tilde{y}, \delta_y)$ can be written as $f_N(n \mid a, x)$ tilted by the likelihood ratio of Y induced by conditioning on (N, A, X) relative to conditioning on (A, X) alone. When $(Y \perp N \mid A, X)$, then $L(\tilde{y}, \delta_y \mid a, x, n)/L(\tilde{y}, \delta_y \mid a, x) = 1$ and $f_{0N}(n \mid a, x, \tilde{y}, \delta_y) = f_N(n \mid a, x)$. When Y and N are associated through unmeasured causes, the tilt deviates from 1 and re-weights values of N that make the observed data $(a, x, \tilde{y}, \delta_y)$ more likely.

We note that expressing $f_{0N}(n \mid a, x, \tilde{y}, \delta_y)$ through the one-way factorization induces dependence between S_{0Y} and f_{0N}, so that these nuisance functions are no longer variation independent.

As a result, not all nuisance-function combinations appearing in the multiple robustness result of Theorem 2 can be achieved independently under this parameterization. Consequently, the set of nuisance configurations yielding consistency may be smaller than that stated in Theorem 2. Nevertheless, the resulting estimator remains doubly robust and retains the same rate-double-robustness property required for asymptotic linearity in Theorem 3.

The remaining challenge is estimation of $S_{0Y}(y \mid a, x, n)$, which is complicated by the fact that both Y and the auxiliary time N may be subject to censoring. However, this can induce selection bias whenever $(N \not\perp Y \mid (A, X))$, since the complete cases need not be representative of the target population. In addition, this approach discards information on the primary outcome (Y), which is particularly inefficient when event rates are low.

The remaining challenge is how to model $S_{0Y}(y \mid a, x, n)$, which is complicated by the fact that the outcome Y and key regressor N may be only partially observed due to censoring. A naive strategy is to restrict estimation to "NCI-complete" cases (i.e., $\Delta_N = 1$). However, this can induce selection bias whenever $N \not\perp Y \mid (A, X)$, since the complete cases need not be representative of the target population. In addition, this approach discards information on the primary outcome (Y), which is particularly inefficient when event rates are low.

To address this issue, we propose a dataset augmentation strategy inspired by the redistribute-to-the-right algorithm[38–40]. The idea is to expand each NCI-censored observation ($\Delta_N = 0$) into a collection of pseudo-observations corresponding to candidate event times exceeding the censoring time. Each pseudo-observation is assigned a weight that redistributes the original observation's mass across the augmented support of N. A toy example of the dataset augmentation procedure is shown in Figure 5.2.

As an initial step, we use working redistribution weights based on the marginal conditional distribution of $N \mid (A, X)$, obtained from the one-way factorization:

$$\hat{w}_j^{(0)} = \hat{f}_N(N_j \mid A, X, N_j > \tilde{N}) = \frac{\hat{f}_N(N_j \mid A, X)}{\hat{S}_N(\tilde{N} \mid A, X)}.$$

A weighted regression of Y on (A, X, N) on the augmented dataset using these weights yields a pilot estimator $\hat{S}_Y^{(0)}$. While this approach recovers marginal information on N, it does not incorporate information from Y, and may therefore be suboptimal when $Y \not\perp N \mid (A, X)$.

To improve efficiency, we refine the weights using an EM-type updating scheme. In the E-step, we update the redistribution weights for NCO-censored observations using the current estimate of the posterior distribution of $N \mid (A, X, \tilde{Y}, \Delta_Y)$:

$$\hat{w}_j^{(m+1)} \propto \hat{L}^{(m)}(\tilde{Y}, \Delta_Y \mid A, X, N_j) \cdot \hat{f}_N(N_j \mid A, X, N_j > \tilde{N}).$$

Where $\hat{L}^{(m)}$ is the estimated likelihood defined above based on the current estimate of the conditional survival function $\hat{S}_Y^{(m)}$ (see 5). In the M-step, we obtain an updated estimator $\hat{S}_{0Y}^{(m+1)}$ by fitting a weighted regression of Y on (A, X, N) using weights $\hat{w}^{(m+1)}$. This procedure is iterated until convergence or until the relative change in the weights satisfies

$$d_{\text{rel}}^{(m)} := \left(\frac{\sum_{i \in \mathcal{I}} (\hat{w}_i^{(m+1)} - \hat{w}_i^{(m)})^2}{\sum_{i \in \mathcal{I}} (\hat{w}_i^{(m)})^2} \right)^{1/2} < \gamma,$$

where $(\mathcal{I})$ indexes observations requiring redistribution.

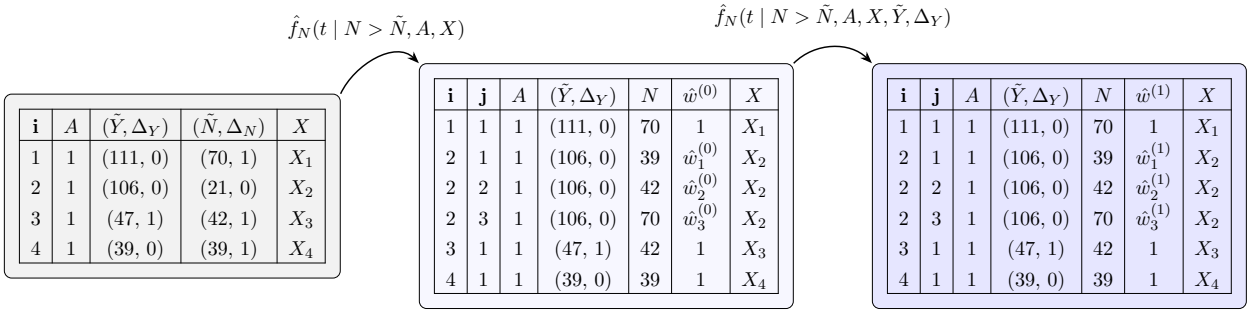

i	A	$(\tilde{Y}, \Delta_Y)$	$(\tilde{N}, \Delta_N)$	X
1	1	(111, 0)	(70, 1)	X_1
2	1	(106, 0)	(21, 0)	X_2
3	1	(47, 1)	(42, 1)	X_3
4	1	(39, 0)	(39, 1)	X_4

i	**j**	A	$(\tilde{Y}, \Delta_Y)$	N	$\hat{w}^{(0)}$	X
1	1	1	(111, 0)	70	1	X_1
2	1	1	(106, 0)	39	$\hat{w}_1^{(0)}$	X_2
2	2	1	(106, 0)	42	$\hat{w}_2^{(0)}$	X_2
2	3	1	(106, 0)	70	$\hat{w}_3^{(0)}$	X_2
3	1	1	(47, 1)	42	1	X_3
4	1	1	(39, 0)	39	1	X_4

i	**j**	A	$(\tilde{Y}, \Delta_Y)$	N	$\hat{w}^{(1)}$	X
1	1	1	(111, 0)	70	1	X_1
2	1	1	(106, 0)	39	$\hat{w}_1^{(1)}$	X_2
2	2	1	(106, 0)	42	$\hat{w}_2^{(1)}$	X_2
2	3	1	(106, 0)	70	$\hat{w}_3^{(1)}$	X_2
3	1	1	(47, 1)	42	1	X_3
4	1	1	(39, 0)	39	1	X_4

Figure 5.2: The data augmentation procedure used to estimate $S_{0Y}(y \mid a, x, n)$.

5.3 Numerical Experiments

We simulated data from randomized trials of $n = \{2400, 4000\}$ participants with $\pi = 0.5$ probability of randomization to vaccine/placebo. We considered trials where the vaccine's prevention efficacy (PE) against the primary infection was 25% (logRR = $\log\{0.75\}$). We assumed that the study administratively censored all infection outcomes at two years, consistent with the design of HVTN 704[41]. We will summarize the performance of estimators based on the average relative bias, median

relative efficiency (compared to unadjusted estimator), and empirical confidence interval coverage of the log risk ratio ($\log\{\mathrm{RR}\}$) of primary infection at the maximum primary event time.

We simulated $\{Y(1), Y(0), C(1), C(0), N, X\}$ from a multivariate normal copula model. We assumed intercorrelation between $Y(1)$, $Y(0)$, and N of $\rho \in \{0, 0.5, 0.8\}$, reflecting cases where the NCI time was unassociated, moderately associated, and highly associated with primary infection times. We assumed the potential dropout times $C(0)$ and $C(1)$ were not intercorrelated with each other nor all other time-to-event variables, thereby satisfying the censoring at random assumption. We considered cases where X was uncorrelated with all time-to-event variables ($\rho_X = 0$). We assumed that $Y(a)$ had a marginal exponential distribution with annual incidence of either 6% in the placebo arm ($\lambda_{Y(a)} = -\log(1 - 0.06 \cdot [1 + a \cdot \mathrm{RR}])$). We assumed that $C(a)$ had a marginal exponential distribution with annual incidence of $\lambda_C = 10\%$ in both arms, reflecting instances of moderate to high censoring. We assumed that N had a marginal exponential distribution with annual incidence of 12% ($\lambda_N = -\log(1 - 0.12)$). We assumed a standard normal marginal distribution of baseline covariate X.

As benchmarks, we considered estimators of $\log \mathrm{RR}(t_0)$ based on unadjusted, naïve Kaplan-Meier curves and estimators that only adjusted for the fully observed, weakly prognostic baseline covariate X[31]. We evaluated two versions of our estimator. The first was an "Oracle" estimator which replaced the unknown nuisance parameters with their known values, which were computed from an external dataset with 10 million observations consistent with the trial's data generating model. The latter version was the proposed cross fitted (CF) estimator, which used a logistic regression model to estimate $\pi_0(a \mid x)$ and cross-fitted SuperLearners to estimate all other nuisances.[31,42] Specifically, we used 10-fold cross-fitted SuperLearner ensembles to estimate $f_N(n \mid a, x)$, and $G_{0Y}(t \mid a, x)$ with Kaplan-Meier, Exponential, Cox, and generalized additive models as base learners. To estimate $S_{0Y}(y \mid a, x, n)$, we used a 10-fold cross-fitted SuperLearner with Cox, generalized additive models, and random forest base learners applied to augmented datasets where censored N were reassigned plausible values of N. We conducted iterative EM-style updates of our estimator until the augmented dataset weight vectors updates were small ($d_{\mathrm{rel}}^{(m)} < 10^{-4}$) or 5 iterations were reached, whichever came first.

The results of the simulation are shown in Table 5.1. All estimators showed low bias ($< 2.5\%$ average relative bias) for the target parameter and achieved near nominal confidence interval cov-

erage. As N became more prognostic for Y, adjusting for N led to greater efficiency gain. When N was moderately prognostic for Y ($\rho = 0.5$), the efficiency gain was $8 - 9\%$. When N was highly prognostic for Y, the efficiency gain was approximately 33% and relative bias decreased. Hence, our proposed estimators exhibited strong performance on par with the benchmark estimators when N is not prognostic, but enhance efficiency when N is prognostic for Y.

Table 5.1: Simulation results at landmark time t_0 based on 1000 replications.

		$\rho = 0.0$			$\rho = 0.5$			$\rho = 0.8$		
n	Method	Rel. Bias (%)	Rel. Eff.	Cov (%)	Rel. Bias (%)	Rel. Eff.	Cov (%)	Rel. Bias (%)	Rel. Eff.	Cov (%)
2000	KM	1.680	1.000	94.9	1.540	1.000	94.1	2.230	1.000	94.4
	X-adj	1.690	1.000	95.1	1.520	1.000	93.9	2.190	1.000	94.3
	NCI-adj (Oracle)	1.710	0.995	95.1	2.080	0.911	93.9	0.528	0.669	94.7
	NCO-adj (CF)	1.760	0.998	94.8	1.920	0.922	94.4	0.425	0.689	94.2
4000	KM	-0.480	1.000	95.4	-0.574	1.000	94.7	0.338	1.000	94.4
	X-Adj	-0.534	1.000	95.6	-0.569	1.000	94.7	0.324	1.000	94.4
	NCO-adj (Oracle)	-0.588	0.997	95.4	-0.721	0.912	94.6	-0.124	0.671	94.2
	NCO-adj (CF)	-0.581	0.998	95.0	-0.715	0.921	93.9	-0.119	0.687	93.6

5.4 Case Study: a randomized trial of a broadly neutralizing antibody for HIV-1 prevention

In the following section, we apply our methods to the HVTN 704 (AMP) study[41], which was a randomized, double-blinded trial evaluating the prevention efficacy of a broadly neutralizing antibody (bNab), VRC01 against HIV-1 infection among 2701 men who have sex with men and transgender women in North America, South America, and Europe. HVTN 704 Participants were randomized in a 1:1:1 ratio to receive low-dose VRC01, high-dose VRC01, and placebo atop standard of prevention. The primary outcome was serologically confirmed HIV-1 infection at 80 weeks of follow up, and HIV-1 testing was conducted every 4 weeks. Participants also underwent sexually transmitted infection testing at enrollment and every 6 months or after presentation of symptoms. Sexually transmitted infection testing included gonorrhea (oropharyngeal, urethral, and rectal), chlamydia (urethral and rectal), and syphilis (serologic)[43].

As a first step, we will evaluate if each STI is a valid NCO for HIV-1 infection. Second, we will explore whether adjustment for time-to-STI can improve the precision of the treatment effect of VRC01 against HIV-1. We focused on two candidate NCOs—time-to-first rectal gonorrhea (RGC)

and syphilis infections—as these were previously suggested as factors independently associated with higher likelihood of HIV-1 [44–47].

As shown in Figure 5.3, we observe that there is no discernible effect of VRC01 administration on incidence of RGC over 72 weeks. Moreover, we observed that binary RGC infection status was associated with higher cumulative incidence of HIV-1 infection over follow-up. In particular, placebo recipients with an incident RGC infection appeared at approximately 50% higher relative risk of acquiring HIV-1 over 80 weeks compared to placebo recipients who avoided RGC infection. These results supported that RGC was a compelling NCO for HIV-1, as it was both unaffected by treatment yet positively associated with HIV-1. Results were similar for syphilis, another STI previously proposed as a marker of HIV-1 exposure.

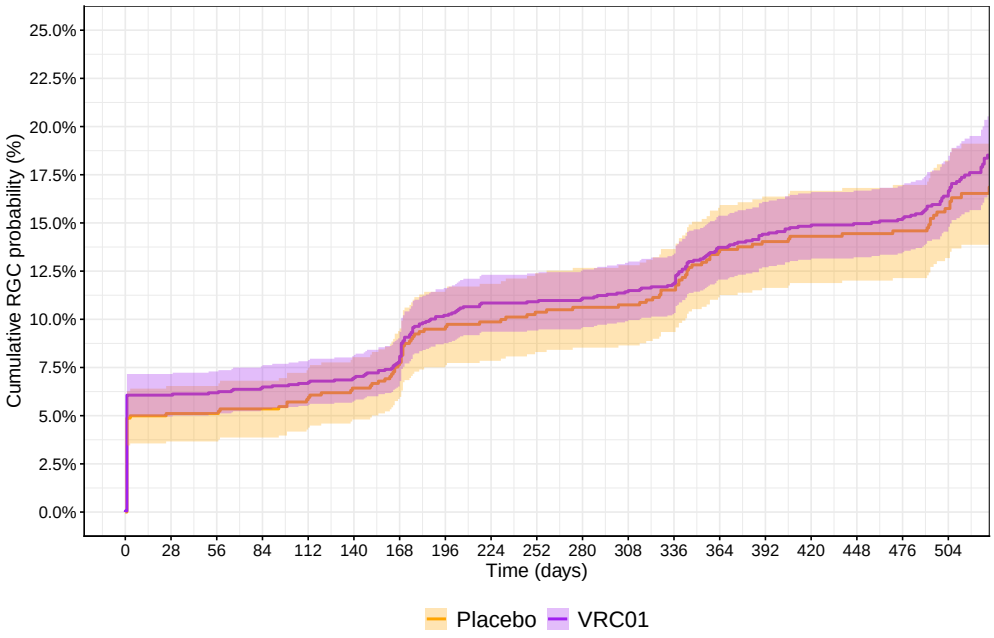

(a) Time-to-first RGC is not influenced by VRC01 in HVTN704.

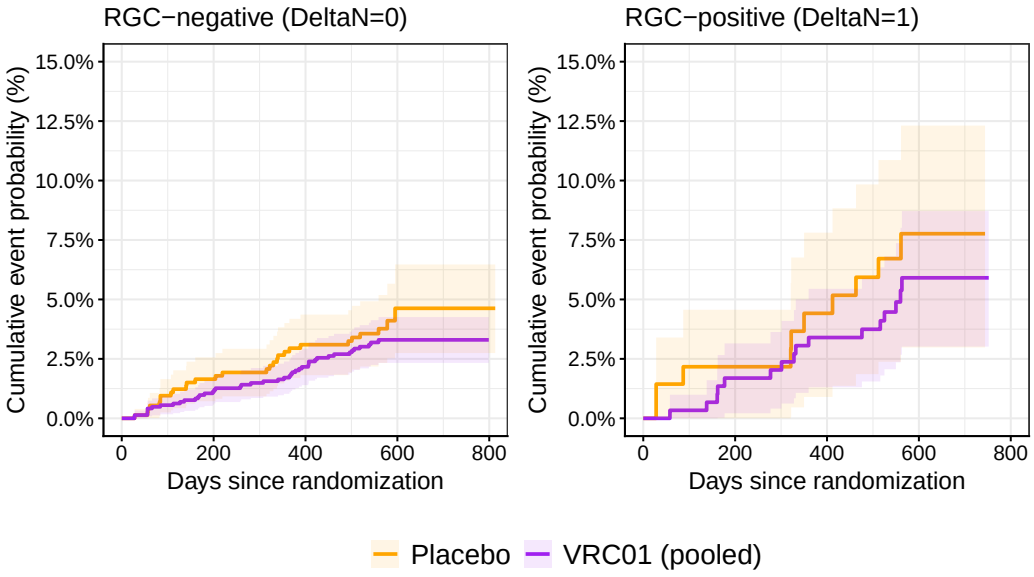

(b) RGC status is positively associated with HIV-1 risk in HVTN704.

Figure 5.3: Time-to-Rectal Gonorrhea (RGC) is a valid NCO for HIV-1 infection in AMP.

We focused on time-to-first rectal gonorrhea (RGC) as the negative control outcome of interest, given it has previously been reported as a marker of HIV exposure. Before adjusting for RGC infection times in the primary analysis, we confirmed that RGC was indeed unaffected by VRC01. To do this, we fit Kaplan-Meier curves stratified by treatment arm to the time-to-RGC outcomes. We administratively censored negative control infection outcomes at each participant's week 72 visit date, which corresponded to their last scheduled STI testing date. The cumulative incidence curves can be found in Figure 5.3. The cumulative incidences of RGC were nearly identical between placebo and VRC01 arms, which supports the biologically plausible assumption that a broadly neutralizing antibody against HIV-1 would not influence the risk of other bacterial STI outcomes. However, we observed that participants in the placebo arm with an RGC episode had roughly 50% higher cumulative incidence of HIV-1 compared to placebo recipients without an RGC episode. These results support that RGC is a compelling choice of NCO for HIV-1, as it was both unaffected by VRC01 administration yet positively associated with HIV-1 outcomes.

Our target estimand was $\log \mathrm{RR}(t_0)$, or the log VRC01:placebo cumulative incidence ratio of HIV-1 at the 80-week visit, where the low and high-dose VRC01 arms were pooled.

We investigated whether adjustment for time-to-incident STIs could improve precision of prevention efficacy estimates of VRC01 in HVTN 704. As a benchmark, we estimated the logRR and PE using (i) an unadjusted estimator based on Kaplan-Meier estimators and (ii) a one-step estimator as described in Westling et al. (2023) which adjusted for a baseline HIV-1 risk score estimated using machine learning[41]. We compared the benchmark estimators to an estimator adjusting for time-to-first rectal gonorrhea (RGC) as the key adjustment variable. We used a logistic regression model to estimate the treatment propensity and an 8-fold cross-fitted SuperLearner to estimate each remaining nuisance. The density $f_N(n \mid a, x)$ was estimated using an ensemble of exponential, Weibull, Cox, and additive Cox models. The survival function $S_{0Y}(y \mid a, x, n)$ was estimated using an ensemble of Cox, additive Cox, and survival random forest. We used the EM algorithm based on data augmentation to estimate S_{0Y} with a maximum of 5 iterations of the EM algorithm per fold. All censoring distributions were estimated using an ensemble of Kaplan-Meier, Exponential, Weibull, Cox, additive Cox, and random forest base learners. All outcomes were summarized at the week 80 visit, consistent with the primary analysis.

All estimators supported that the pooled VRC01 group had moderately lower incidence of HIV

compared to the placebo arm. While adjusting for the baseline risk score (X) only achieved a modest 2.5% reduction in the estimated variance of the target estimand, our proposed estimator which adjusted for time-to-first RGC infection exhibited a 26.9% variance reduction relative to the unadjusted analysis. Similarly, adjusting for time-to-first syphilis infection reduced the estimated variance by 26.5%.

Estimator	logRR	Est. Var	Relative Eff.	PE (95% CI)
Kaplan-Meier	-0.309	4.60e-2	1	$26.6\%(-11.7\%, 51.8\%)$
X-adjusted	-0.294	4.49e-2	0.975	$25.4\%(-12.9\%, 50.8\%)$
RGC-adjusted	-0.248	3.36e-2	0.731	$22.0\%(-11.7\%, 45.6\%)$
Syph-adjusted	-0.261	3.38e-2	0.735	$23.0\%(-10.5\%, 46.3\%)$

Table 5.2: Comparison of different estimators of VRC01 prevention efficacy in HVTN 704. Estimates are reported at 728 days (Week 104).

5.5 Discussion

Herein, we develop a framework to enhance the efficiency of randomized trials of infectious disease prevention methods by adjusting for potentially right-censored negative control outcome (NCO) event times. Using the theory of coarsened data influence functions, we derived the efficient influence function (EIF) for the treatment-specific survival function of the primary infection outcome when the NCO event time was also right-censored. We proposed a cross-fitted, one-step estimator which is asymptotically linear and multiply robust under conditions that encourage the use of flexible, data-adaptive nuisance estimators. We proposed an iterative EM algorithm based on data augmentation and redistribute-to-the-right to estimate one of the key nuisance parameters. In simulations, we demonstrate our proposed estimator is consistent and leads to well-calibrated inference and tangible efficiency gains relative to several benchmark estimators. Finally, we applied our proposed method to the HVTN704 study which evaluated the efficacy of a broadly neutralizing antibody for preventing of HIV-1 acquisition in a cohort of 2701 men who have sex with men and transgender women in the Americas and Europe. While adjusting for baseline covariates led to very modest precision gains, adjusting for two different time-to-sexually transmitted infections resulted in efficiency gains of approximately 27% relative to unadjusted benchmarks.

Our work contributes to the growing statistical literature on regression with censored covariates,[48–58] particularly those that view the problem from the lens of semiparametric theory in coarsened data models.[32,33,59–63] Our work is the first to our consider the problem of right-censored covariate and outcomes in the context of infectious disease prevention trials. Our approach extends that of Etievant et al. (2023) which proposed adjusting for binary indicators of observed "vaccine untargeted" HPV infections in HPV vaccine trials. However, our approach more naturally accommodates the typical time-to-event nature of infection endpoints in prevention trials, avoids information loss by binarization, and eschews the need for a completely random censoring assumption. We highlight that NCO event times may be particularly useful adjustment variables in trials with longer follow up periods and in trials that employ frequent and/or symptom-prompted testing for negative control infections. Studies with very short follow up durations or very infrequent negative control infection testing may admit lower efficiency gains when adjusting for NCO event times. While our work was motivated by randomized infectious disease prevention trials with time-to-event endpoints, in principle, our method could be applied to a randomized trial or observational study where the primary outcome is time-to-event and there exists a compelling negative control time-to-event outcome that may be right-censored.

We emphasize that while the negative control assumption is strong, it may be biologically plausible for specific treatments and can be tested with reasonably high power in randomized prevention trials. Typically variables measured post-baseline are excluded from adjustment due to the potential for selection bias, although this risk is absent when the post-baseline variable is unaffected by treatment[23]. In our data application to HVTN704, we observed no discernible effect of VRC01 on cumulative incidence of rectal gonorrhea, which supports the validity of the negative control assumption. Moreover, it is highly implausible that a passively administered antibody against the HIV-1 virus would influence the likelihood of acquiring a bacterial STI, due to the antigenic specificity of the adaptive immune response. We believe the negative control assumption is most credible for antigenically specific interventions such as vaccines and passively administered antibodies. We cannot confirm with certainty if our assumption holds, empirical testing coupled with scientific knowledge provide strong justification for this assumption in many prevention studies.

In future work, it may be of interest to extend our framework to accommodate multivariate and/or recurrent right-censored negative control infection endpoints. In our developments, we as-

sume a single, time-to-first negative control infection is measured in parallel to the outcome of interest. However, in HVTN704, various STIs such as gonorrhea (oropharyngeal, urethral, and rectal), chlamydia (urethral and rectal), and syphilis (treponemal and non-treponemal serology) were measured throughout follow-up. Each STI outcome is a noisy proxy of factors that influence an individual's HIV exposure, so including multiple STI outcomes in the analysis may offer additional efficiency gain. To extend the work herein to accommodate multivariate, recurrent negative control infections, we could employ conditional imputation of multivariate right-censored survival outcomes from sequences of conditionally specified survival models[40,64,65]. Moreover, many STIs do not produce sterilizing immunity, so persistent/recurrent infections can occur. Chronic STIs have been found to transiently increase risk of HIV-1 transmission and acquisition[18,19]. Accounting for persistent or recurrent STIs may provide additional explanatory information about HIV risk exposure beyond the time-to-first STI, which may generate additional gains in precision and power. Moreover, adjusting for quantitative readouts of STI tests (e.g., viral/bacterial load) could offer additional precision gains by describing the severity of the negative control infection endpoint.

Bibliography

[1] Li Yang and Anastasios A. Tsiatis. Efficiency study of estimators for a treatment effect in a pretest-posttest trial. *The American Statistician*, 55:314–321, 2001.

[2] Anastasios A. Tsiatis, Marie Davidian, Min Zhang, and Xiaomin Lu. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Statistics in medicine*, 27(23):4658–4677, October 2008. ISSN 0277-6715. doi: 10.1002/sim.3113. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2562926/`.

[3] Min Zhang, Anastasios A. Tsiatis, and Marie Davidian. Improving Efficiency of Inferences in Randomized Clinical Trials Using Auxiliary Covariates. *Biometrics*, 64(3):707–715, 2008. ISSN 1541-0420. doi: 10.1111/j.1541-0420.2007.00976.x. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1541-0420.2007.00976.x`. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1541-0420.2007.00976.x.

[4] Winston Lin. Agnostic notes on regression adjustments to experimental data: Reexamining freedman's critique. *Annals of Applied Statistics*, 7:295–318, 2013.

[5] Ting Ye, Jun Shao, Yanyao Yi, and Qingquan Zhao. Toward better practice of covariate adjustment in analyzing randomized clinical trials. *Journal of the American Statistical Association*, 2022.

[6] Kelly Van Lancker, Frank Bretz, and Oliver Dukes. Covariate adjustment in randomized controlled trials: General concepts and practical considerations. *Clinical Trials*, 2024.

[7] C.P. Farrington, M.N. Kanaan, and N.J. Gay. Estimation of the basic reproduction number for infectious diseases from age-stratified serological survey data. *Applied Statistics*, 50:251–292, 2001.

[8] C Paddy Farrington, Heather J Whitaker, Steffen Unkel, and Richard Pebody. Correlated infections: quantifying individual heterogeneity in the spread of infectious diseases. *Am. J. Epidemiol.*, 177(5):474–486, March 2013.

[9] Kerstin E.E. Schroder, Michael P. Carey, and Peter A. Vanable. Methodological challenges in research on sexual risk behavior: Ii. accuracy of self-reports. *Annals of Behavioral Medicine*, 26:104—-123, 2003.

[10] C. Paddy Farrington and Heather J Whitaker. Contact Surface Models for Infectious Diseases: Estimation From Serologic Survey Data. *Journal of the American Statistical Association*, 100 (470):370–379, June 2005. ISSN 0162-1459, 1537-274X. doi: 10.1198/016214504000001754. URL `http://www.tandfonline.com/doi/abs/10.1198/016214504000001754`.

[11] C. Paddy Farrington, Heather J. Whitaker, Steffen Unkel, and Richard Pebody. Correlated Infections: Quantifying Individual Heterogeneity in the Spread of Infectious Diseases. *American Journal of Epidemiology*, 177(5):474–486, March 2013. ISSN 1476-6256, 0002-9262. doi: 10.1093/aje/kws260. URL `https://academic.oup.com/aje/article-lookup/doi/10.1093/aje/kws260`.

[12] Katharine J Looker, Jocelyn A R Elmes, Sami L Gottlieb, Joshua T Schiffer, Peter Vickerman, Katherine M E Turner, and Marie-Claude Boily. Effect of HSV-2 infection on subsequent HIV acquisition: an updated systematic review and meta-analysis. *Lancet Infect. Dis.*, 17(12): 1303–1316, December 2017.

[13] Chris Kenyon. Strong associations between national prevalence of various STIs suggests sexual network connectivity is a common underpinning risk factor. *BMC Infect. Dis.*, 17(1):682, October 2017.

[14] Ross A Baiers, Daniel T Ryan, Antonia Clifford, Erik Munson, Richard D'Aquila, Michael E Newcomb, and Brian Mustanski. Asymptomatic rectal bacterial pathogens show large prospective relationships with HIV incidence in a cohort of young sexual and gender minorities: Implications for STI screening and HIV prevention. *Open Forum Infect. Dis.*, 11(8):ofae444, August 2024.

[15] Mark Richard Stenger, Preeti Pathela, Christina Schumacher, Nicole Burghardt, Rachel Amiya, Robbie Madera, Trang Q Nguyen, and Elizabeth Torrone. Trends in HIV prevalence by self-report among MSM diagnosed and reported with gonorrhea in six united states jurisdictions from 2010 to 2019. *AIDS*, 35(15):2523–2530, December 2021.

[16] Heather N Grome, Peter F Rebeiro, Meredith Brantley, Dyanne Herrera-Vasquez, Samantha A Mathieson, and April C Pettit. Risk of HIV diagnosis following bacterial sexually transmitted infections in tennessee, 2013-2017. *Sex. Transm. Dis.*, 48(11):873–880, November 2021.

[17] Charu Mullick and Jeffrey Murray. Correlations between human immunodeficiency virus (HIV) infection and rectal gonorrhea incidence in men who have sex with men: Implications for future HIV preexposure prophylaxis trials. *J. Infect. Dis.*, 221(2):214–217, January 2020.

[18] Helen M Chun, Robert J Carpenter, Grace E Macalino, and Nancy F Crum-Cianflone. The role of sexually transmitted infections in HIV-1 progression: A comprehensive review of the literature. *J. Sex. Transm. Dis.*, 2013:176459, June 2013.

[19] Myron S Cohen, Olivia D Council, and Jane S Chen. Sexually transmitted infections and HIV in the era of antiretroviral treatment and prevention: the biologic basis for epidemiologic synergy. *J. Int. AIDS Soc.*, 22 Suppl 6(S6):e25355, August 2019.

[20] David Kim, James Quinn, Benjamin Pinsky, Nigam H Shah, and Ian Brown. Rates of co-infection between SARS-CoV-2 and other respiratory pathogens. *JAMA*, 323(20):2085–2086, May 2020.

[21] Ethan Ashby, Holly Janes, Dean Follmann, Peter B. Gilbert, Honghong Zhou, Xiaowei Wang, Bethany Girard, James B. Kublin, Lawrence Corey, Kathleen M. Neuzil, Lindsey R. Baden, Hana M. El Sahly, and Bo Zhang. Validating and leveraging non-sars-cov-2 respiratory infection as a negative control outcome in a phase 3 covid-19 vaccine trial with extended observational follow-up. *The American Journal of Epidemiology*, 2025. https://doi.org/10.1093/aje/kwaf176.

[22] Miguel A Hernán and James M Robins. *Causal Inference: What If*. Chapman and Hall/CRC, 2020.

[23] Paul R. Rosenbaum. The consequences of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society: Series B*, 147, 1984.

[24] Marc Lipsitch, Eric Tchetgen Tchetgen, and Ted Cohen. Negative Controls: A Tool for Detecting Confounding and Bias in Observational Studies. *Epidemiology (Cambridge, Mass.)*, 21(3):383–388, May 2010. ISSN 1044-3983. doi: 10.1097/EDE.0b013e3181d61eeb. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3053408/`.

[25] Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. A Selective Review of Negative Control Methods in Epidemiology. *Current epidemiology reports*, 7(4):190–202, December 2020. ISSN 2196-2995. doi: 10.1007/s40471-020-00243-4. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8118596/`.

[26] Lola Etievant, Joshua M. Sampson, and Mitchell H. Gail. Increasing efficiency and reducing bias when assessing hpv vaccination efficacy by using nontargeted hpv strains. *Biometrics*, 79 (2):1534—-1545, 2023.

[27] Jerzy Neyman. Sur les applications de la theorie des probabilites aux experiences agricoles: Essai des principes. *Statistical Science*, 5, 1921.

[28] Donald B Rubin. Causal inference using potential outcomes. *Journal of the American Statistical Association*, 100, 2005.

[29] Michael P Fay, Dean Follmann, Bruce J Swihart, and Lauren E Dang. Vaccine efficacy estimands implied by common estimators used in individual randomized field trials. January 2026.

[30] Donald B Rubin. Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, 6, 1978.

[31] Ted Westling, Alex Luedtke, Peter B. Gilbert, and Marco Carone. Inference for treatment-specific survival curves using machine learning. *Journal of the American Statistical Association*, 2023.

[32] Anastasios A Tsiatis. *Semiparametric Theory and Missing Data*. Springer Series in Statistics. Springer, New York, NY, 2006 edition, June 2006.

[33] Peter J. Bickel, Chris A.J. Klaassen, Ya'acov Ritov, and Jon A. Wellner. *Efficient and adaptive estimation for semiparametric models*. Johns Hopkins University Press Baltimore, 1993.

[34] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21, 2018.

[35] Mark J van der Laan, Eric C Polley, and Alan E Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6:number 1, 2007.

[36] Barry C Arnold and S. James Press. Compatible conditional distributions. *J. Am. Stat. Assoc.*, 84(405):152, March 1989.

[37] Hua Yun Chen. Compatibility of conditionally specified models. *Stat. Probab. Lett.*, 80(7-8): 670–677, April 2010.

[38] Bradley Efron. he two sample problem with censored data. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 4, pages 831–853. University of California Press, Berkeley, CA, 1967.

[39] Martin A Tanner and Wing Hung Wong. The calculation of posterior distributions by data augmentation. *J. Am. Stat. Assoc.*, 82(398):528, June 1987.

[40] Jeremy MG Taylor, Susan Murray, and Chiu-Hsieh Hsu. Survival estimation and testing via multiple imputation. *Statistics and Probability Letters*, 58:221–232, 2002.

[41] Lawrence Corey, Peter B. Gilbert, Michal Juraska, David C. Montefiori, Lynn Morris, Shelly T. Karuna, Srilatha Edupuganti, Nyaradzo M. Mgodi, Allan C. deCamp, Erika Rudnicki, Yunda Huang, Pedro Gonzales, Robinson Cabello, Catherine Orrell, Javier R. Lama, Fatima Laher, Erica M. Lazarus, Jorge Sanchez, Ian Frank, Juan Hinojosa, Magdalena E. Sobieszczyk, Kyle E. Marshall, Pamela G. Mukwekwerere, Joseph Makhema, Lindsey R. Baden, James I. Mullins, Carolyn Williamson, John Hural, M. Juliana McElrath, Carter Bentley, Simbarashe Takuva, Margarita M. Gomez Lorenzo, David N. Burns, Nicole Espy, April K. Randhawa, Nidhi Kochar, Estelle Piwowar-Manning, Deborah J. Donnell, Nirupama Sista, Philip Andrew, James G. Kublin, Glenda Gray, Julie E. Ledgerwood, John R. Mascola, and Myron S. Cohen. Two Randomized Trials of Neutralizing Antibodies to Prevent HIV-1 Acquisition. *New England Journal of Medicine*, 384(11):1003–1014, March 2021. ISSN 0028-4793, 1533-4406. doi: 10.1056/NEJMoa2031738. URL `http://www.nejm.org/doi/10.1056/NEJMoa2031738`.

[42] Mark J. van der Laan, E. Polley, and A. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6, 2007.

[43] Srilatha Edupuganti, Nyaradzo Mgodi, Shelly T. Karuna, Philip Andrew, Erika Rudnicki, Nidhi Kochar, Allan deCamp, Robert De La Grecca, Maija Anderson, Carissa Karg, India Tindale, Elizabeth Greene, Gail B. Broder, Jonathan Lucas, John Hural, Jorge A. Gallardo-Cartagena, Pedro Gonzales, Ian Frank, Magdalena Sobieszczyk, Margarita M. Gomez Lorenzo, David Burns, Peter L. Anderson, Maurine D. Miner, Julie Ledgerwood, John R. Mascola, Peter B. Gilbert, Myron S. Cohen, Lawrence Corey, and the HVTN 704/HPTN 085 study group. Feasibility and Successful Enrollment in a Proof-of-Concept HIV Prevention Trial of VRC01, a Broadly Neutralizing HIV-1 Monoclonal Antibody. *JAIDS Journal of Acquired Immune Deficiency Syndromes*, 87(1):671–679, May 2021. ISSN 1525-4135. doi: 10.1097/QAI.0000000000002639. URL `https://journals.lww.com/10.1097/QAI.0000000000002639`.

[44] K J Craib, D R Meddings, S A Strathdee, R S Hogg, J S Montaner, M V O'Shaughnessy, and

M T Schechter. Rectal gonorrhoea as an independent risk factor for HIV infection in a cohort of homosexual men. *Genitourin. Med.*, 71(3):150–154, June 1995.

[45] Adam S Vaughan, Colleen F Kelley, Nicole Luisi, Carlos del Rio, Patrick S Sullivan, and Eli S Rosenberg. An application of propensity score weighting to quantify the causal effect of rectal sexually transmitted infections on incident HIV among men who have sex with men. *BMC Med. Res. Methodol.*, 15(1):25, March 2015.

[46] Lindley A Barbee, Christine M Khosropour, Julia C Dombrowksi, and Matthew R Golden. New human immunodeficiency virus diagnosis independently associated with rectal gonorrhea and chlamydia in men who have sex with men. *Sex. Transm. Dis.*, 44(7):385–389, July 2017.

[47] Meng Yin Wu, Hui Zi Gong, Kui Ru Hu, He-Yi Zheng, Xia Wan, and Jun Li. Effect of syphilis infection on HIV acquisition: a systematic review and meta-analysis. *Sex. Transm. Infect.*, 97 (7):525–533, November 2021.

[48] Ryan C May, Joseph G Ibrahim, and Haitao Chu. Maximum likelihood estimation in generalized linear models with multiple covariates subject to detection limits. *Stat. Med.*, 30(20): 2551–2561, September 2011.

[49] Huixia Judy Wang and Xingdong Feng. Multiple imputation for M -regression with censored covariates. *J. Am. Stat. Assoc.*, 107(497):194–204, March 2012.

[50] Paul W Bernhardt, Huixia J Wang, and Daowen Zhang. Statistical methods for generalized linear models with covariates subject to detection limits. *Stat. Biosci.*, 7(1):68–89, May 2015.

[51] Folefac D Atem, Jing Qian, Jacqueline E Maye, Keith A Johnson, and Rebecca A Betensky. Multiple imputation of a randomly censored covariate improves logistic regression analysis. *J. Appl. Stat.*, 43(15):2886–2896, March 2016.

[52] Folefac D. Atem, Jing Qian, Jacqueline E. Maye, Keith A. Johnson, and Rebecca A. Betensky. Linear regression with a randomly censored covariate: application to an alzheimer's study. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 66, 2017.

[53] Ying Ding, Shengchun Kong, Shan Kang, and Wei Chen. A semiparametric imputation approach for regression with censored covariate with application to an AMD progression study. *Stat. Med.*, 37(23):3293–3308, October 2018.

[54] Folefac D. Atem, Roland A. Matsouaka, and Vincent E. Zimmern. Cox regression model with randomly censored covariates. *Biometrical Journal*, 61, 2019.

[55] Roland A. Matsouaka and Folefac D. Atem. Regression with a right-censored predictor using inverse probability weighting methods. *Statistics in Medicine*, 39, 2020.

[56] Ling-Wan Chen, Jason P Fine, Eric Bair, Victor S Ritter, Thomas F McElrath, David E Cantonwine, John D Meeker, Kelly K Ferguson, and Shanshan Zhao. Semiparametric analysis of a generalized linear model with multiple covariates subject to detection limits. *Stat. Med.*, 41(24):4791–4808, October 2022.

[57] Marissa C Ashner and Tanya P Garcia. Understanding the implications of a complete case analysis for regression models with a right-censored covariate. *Am. Stat.*, 78(3):335–344, 2024.

[58] Sarah C Lotspeich, Marissa C Ashner, Jesus E Vazquez, Brian D Richardson, Kyle F Grosser, Benjamin E Bodek, and Tanya P Garcia. Making sense of censored covariates: Statistical methods for studies of huntington's disease. *Annu. Rev. Stat. Appl.*, 11(1):255–277, April 2024.

[59] James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *J. Am. Stat. Assoc.*, 89(427):846–866, September 1994.

[60] Seong-Ho Lee, Brian D Richardson, Yanyuan Ma, Karen S Marder, and Tanya P Garcia. Robust and efficient estimation in the presence of a randomly censored covariate. 2024.

[61] Jesus E Vazquez, Yanyuan Ma, Karen Marder, and Tanya P Garcia. Robust estimation under outcome dependent right censoring in huntington disease: Estimators for low and high censoring rates. November 2025.

[62] Zhewei Zhang, Yanyuan Ma, Karen Marder, and Tanya P Garcia. Super doubly robust and efficient estimator for informative covariate censoring. November 2025.

[63] Jesus E Vazquez, Marissa C Ashner, Yanyuan Ma, Karen Marder, and Tanya P Garcia. Establishing the parallels and differences between right-censored and missing covariates. September 2024.

[64] S Van Buuren, J P L Brand, C G M Groothuis-Oudshoorn, and D B Rubin. Fully conditional specification in multivariate imputation. *J. Stat. Comput. Simul.*, 76(12):1049–1064, December 2006.

[65] Jonathan W Bartlett, Shaun R Seaman, Ian R White, James R Carpenter, and Alzheimer's Disease Neuroimaging Initiative*. Multiple imputation of covariates by fully conditional specification: Accommodating the substantive model. *Stat. Methods Med. Res.*, 24(4):462–487, August 2015.

Chapter 6
DISCUSSION AND FUTURE DIRECTIONS

6.1 Discussion

In this dissertation, we address two challenges at the frontier vaccine effectiveness evaluation: (i) how to more reliably infer waning vaccine effectiveness over time and (ii) enable efficient inference on vaccine efficacy in randomized trials. These are important objectives, as understanding waning vaccine effectiveness is critical for guiding vaccine science and public health decision-making, such as when to offer updated booster vaccines to the public. Obtaining more precise, efficient inference on vaccine efficacy in randomized trials can help boost the power of vaccine studies, improve confidence in go/no-go decisions during vaccine development pipeline, and may more efficient vaccine trials.

Our approaches leverage negative control outcomes, or outcomes causally unaffected by the vaccine of interest, to address the scientific problems described in this dissertation. We emphasize that vaccine studies are a compelling setting to study negative controls, because candidate negative controls can be readily proposed and justified on the basis of scientific knowledge and empirical evidence. Because vaccines are antigenically specific interventions, candidate negative control outcomes may be readily conceptualized by vaccine scientists and can be measured with low additional effort. Moreover, as illustrated in Chapter 2, randomized vaccine trials permit a unique opportunity to validate and compare candidate negative control outcomes, which can inform the conduct of secondary non-randomized analyses and future observational studies. Lastly, we highlight that understanding many important aspects of vaccine effectiveness — such as effectiveness in special populations of interest, against novel variants, of different vaccination schedules, and on secondary outcomes like transmission and disease progression — may rely on observational data sources. Measuring and utilizing valid negative controls in such analyses can provide a useful indicator of whether assumptions are met or whether bias is contaminating the analysis.

Some of the methods described herein may be applicable beyond vaccine studies. For example, the methods described in Chapters 2 and 3 represent novel approaches to detect and reduce bias in time-varying hazard ratios. In principle, these approaches could be applied to studies of other

kinds of biomedical interventions, as long as a valid negative control is measured in parallel to the outcome of interest. Moreover, the methods described in Chapters 4 and 5 could be used to improve the efficiency in randomized trials of other kinds of interventions, provided there are continuous/time-to-event primary and negative control outcomes available to the analyst.

Broadly, we believe that this dissertation underscores the value of measuring and creatively utilizing negative control outcomes in vaccine studies. In particular, in large observational studies or late-phase randomized trials, negative control infection outcome may be collected with low additional effort. Even though such outcomes are typically deemed "vaccine-irrelevant" and discarded from primary analyses, they may have hidden utility in probing secondary analyses for bias, augmenting the primary analysis, and guiding future observational studies. We hope that this dissertation supports the collection of negative controls and inspires additional statistical methods development.

6.2 Future Directions

Herein, we describe some areas for future research related to the primary aims of this dissertation.

6.2.1 Inference on waning VE in Test Negative Designs

In Chapter 3, we explored the potential of using negative control infections to eliminate confounding and depletion of susceptibles bias from hazard-based, time-varying vaccine effectiveness estimates. Using a fairly general set of assumptions, we showed that the time-varying VE could be identified using a logistic regression model for the odds of experiencing a vaccine-preventable infection ($J = 1$) relative to a vaccine-irrelevant infection ($J = 0$) conditional on some infection occurring on date t.

$$\log\left\{\frac{P(J=1|V=v,T=t)}{P(J=0|V=v,T=t)}\right\} = I(v \leq t)f(t-v) + \alpha_0(t)$$

Where $\text{VE}_1(t-v) = 1 - \exp\{f(t-v)\}$ quantifies the time-varying VE.

This approach is well-suited for use in test negative designs (TNDs).[1] TND studies enroll patients seeking healthcare in response to disease symptoms, and participants are tested for the target pathogen and their vaccination status is ascertained. Participants testing positive for the target pathogens are designated "cases" while those negative for the target pathogen are designated "con-

trols". If a TND study also collects a participant's vaccination date, then the estimator above can be directly applied to infer time-varying VE. Moreover, since TNDs only recruit participants with disease symptoms and the logistic regression model above conditions on infection/disease, our method could pair very efficiently with TNDs. It will be important to formalize assumptions under which the above model identifies waning VE in TNDs, akin to work by Boyer et al. (2026) for time-invariant VE estimates.

In Chapter 3, we restricted focus to inference in the semiparametric working model where we assumed $f(t-v) = \beta_0^T \boldsymbol{\psi}(t-v)$ for a known d-dimensional basis and unknown $\beta_0 \in \mathbb{R}^d$. This was done for the purposes of obtaining consistent and asymptotically normal estimators of time-varying VE. As alluded to in Chapter 3, we may wish to impose the more substantive constraint that $f(t-v)$ be monotone non-decreasing in its input, thereby ensuring that vaccine efficacy is monotone non-increasing in time-since-vaccination. This respects the scientific belief that after a brief ramp-up period, a vaccine's effectiveness should either remain constant or wane over time, but should not increase over time. To address this problem, particularly while leaving the nuisance parameter $\alpha_0(t)$ unspecified, we likely need to exploit techniques developed for inference on function-valued parameters under monotone shape constraints.[3,4]

6.2.2 Methods applications to CASCADIA

In future work, we may apply the methods developed in Chapters 2 and 3 to analyze time-varying vaccine effectiveness in the CASCADIA study.[5] CASCADIA is a 4-year community-based prospective study of SARS-CoV-2 VE that enrolled children 6 months to 17 years old and adults aged 18–49 years. CASCADIA participants were instructed to complete weekly self-collection of anterior nasal swabs and symptom questionnaires. Swabs are tested for SARS-CoV-2 and other respiratory pathogens (Influenza A and B, RSV A and B, and others) by reverse transcription-PCR.

Prior analysis of data pooled from CASCADIA and other prospective cohort studies examined the efficacy of monovalent and bivalent COVID-19 vaccines in children,[6] infants,[7] and adult populations.[8] Some of these analyses reported VE estimates over discrete non-overlapping intervals in days-since-vaccination. However, these analyses did not account for potential unmeasured confounding and depletion of susceptibles bias, and therefore are vulnerable to bias.

To address this, we could apply approaches described in Chapters 2 and 3 to detect and adjust for unmeasured confounding and selection bias using the other "vaccine-irrelevant" respiratory pathogens that were measured via the reverse transcription-PCR. One key consideration is that voluntary COVID-19 vaccination may be associated with voluntary receipt of other vaccines (e.g., influenza, RSV, etc). Correlated vaccine uptake could lead to lower incidence of the negative control infections, which could lead to an incorrect assumption that COVID-19-vaccine recipients are less exposed to respiratory pathogens. Hence, any analysis should either adjust for vaccination status against other vaccine-preventable pathogens, or focus on other pathogens (human rhinovirus) that correlated vaccine uptake is unlikely to influence.

6.2.3 Exploiting natural exposure gradients to infer waning VE in traveler cohorts

We briefly summarize the work on inferring waning influenza VE by Lipsitch et al., Ray et al. (2019, 2020), which serves as the seminal motivation for this idea. The authors recognized that traditional, hazard-based estimates of time-varying VE are not reliable indicators of waning VE due to depletion of susceptibles bias (see Chapter 1.2 for extended discussion of this bias and how it arises). However, the authors noted that specific pathogens like influenza mainly circulate only during a part of the year (mid-December to March). The authors formulated a hypothetical randomized trial comparing infection/disease rates between cohorts vaccinated at different points before the start of the flu season. This essentially creates to a randomized comparison between "early" and "late" vaccination, which permits testing of VE waning.[11] This comparison is also robust to depletion of susceptibles bias, as the predictable temporal variation in influenza exposure enables a comparison of "early" and "late" vaccination where exposure (and necessarily infection) is naturally deferred until after all participants are vaccinated.[12]

We can naturally extend these ideas to any pathogen that exhibits predictable variation in exposure along a readily observable axis. We term these predictable variations as "exposure gradients". While some pathogens like influenza exhibit seasonal exposure gradients, many pathogens of interest to global health exhibit dramatic exposure gradients in space. For example, malaria exhibits marked variation in space by virtue of the geographic distribution of the *Anopheles* mosquito, the carrier of the pathogen parasite *Plasmodium* (Figure 6.1). Other pathogens such as dengue,

chikungunya, Zika, and yellow fever all exhibit marked variation in incidence over space that tracks the geographic distribution of their vectors.[13]

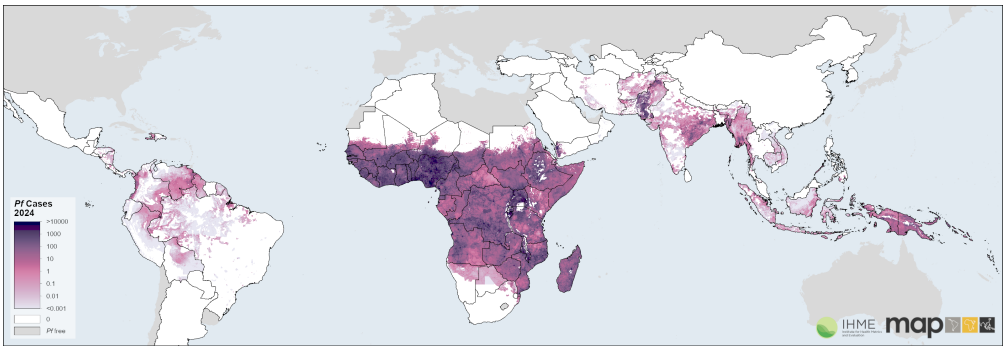

Figure 6.1: Example of spatial exposure gradient for *Plasmodium falciparum*, one of the causative agents for malaria.[14]

Suppose we wish to infer if a malaria vaccine's effectiveness wanes over time. To assess this, we envision a hypothetical randomized study that enrolls participant planning on traveling from a "low incidence" region (e.g., South Africa) to a "high incidence" region (e.g., Nigeria) in the near future. The trial would randomize participants to vaccination either several months (early) or several weeks (late) before their anticipated travel date and compare the incidence of malaria between these groups during travel. Similar to the influenza case, this study design is robust to depletion of susceptibles as prior to the date of travel, the participants cannot be exposed to malaria and therefore depletion of susceptibles cannot manifest. Whereas the start of the flu season is the "exposure switch" for influenza, the analog for malaria is travel from the low to high-incidence region.

This intuition can be formalized by following the arguments of Lipsitch et al. (2019) and simulation-based stress-testing. An outstanding question is what is the best strategy for reliably inferring waning vaccine efficacy when only observational data are available. One approach could assume participants are exchangeable in their infection outcomes after adjusting for known confounders, although this approach could lead to bias when some confounders are unmeasured. A potentially appealing alternative could be to use geographic pair matching,[15] with the hope that controlling for geography controls for unmeasured factors influencing exposure, which is especially

appealing for vector-borne pathogens. In future work, we may consider other naturally occurring exposure gradients in vaccine-preventable infections, such as those that arise as infants and young children entering group childcare settings.[16]

6.2.4 Adjusting for multivariate negative control infection times to improve the precision of late-phase vaccine efficacy trials

In Chapter 5, we propose an estimator of vaccine efficacy in a randomized trial which adjusts for the time-to-negative control infection to improve precision, power, and efficiency. Our work was motivated by the HVTN704 (AMP) study, a study of a broadly neutralizing antibody to prevent HIV-1 acquisition in men who have sex with men.[17] Throughout the study, scheduled and symptom-triggered testing for gonorrhea (oropharyngeal, urethral, and rectal), chlamydia (urethral and rectal), and syphilis (treponemal and non-treponemal serology) was conducted.[18?] It seems reasonable that each of the bacterial STIs would be unaffected by broadly neutralizing antibody administration, but each is a noisy proxy for HIV-1 exposure.

Hence, it would be interesting to extend our approach developed in Chapter 5 to adjust for multivariate NCI times, $\boldsymbol{N} = (N^{(1)}, \ldots, N^{(k)})$ which would be useful when a randomized trial tracks several negative control infections in parallel. Each component contributes a separate coarsening of the full data via its own censoring time $C^{(k)}$. The full-data EIF generalizes directly, but the observed-data projection requires modeling the joint survival of the censoring times.

$$
\begin{aligned}
G_{0N}(\boldsymbol{u} \mid a, x, \tilde{y}, \delta_y) :&= P_0(C^{(1)} \geq u_1, \ldots, C^{(K)} \geq u_K \mid a, x, \tilde{y}, \delta_y) \\
&= \prod_{k=1}^{K} \exp\left\{ - \int_0^{u_k} \lambda_{C^{(k)}}(s \mid a, x, \tilde{y}, \delta_Y, C^{(1)} \geq u_1, \ldots, C^{(k-1)} \geq u_{k-1}) dx \right\}
\end{aligned}
$$

Where the second equality sequentially models the conditional hazards of the censoring variables under any chosen ordering of the components. In the case where all outcomes are subject to the same censoring time, a single censoring model fit to all endpoints will suffice.

The augmentation function will take the form

$$
Q_0(\boldsymbol{u} \mid \cdot) := \mathbb{E}[\phi^F_{a_0,t_0}(\ldots, N^{(1)}, \ldots, N^{(K)}) \mid A = a, X = x, \tilde{Y} = \tilde{y}, \Delta_Y = \delta_y, N^{(1)} \geq u_1, \ldots, N^{(K)} \geq u_K].
$$

The leading term fo the multivariate EIF will be

$$\frac{\prod_{k=1}^{K} \delta_n^{(k)}}{G_{0N}(\tilde{\boldsymbol{n}} \mid \cdot)} \phi^F_{a_0,t_0}(\tilde{y}, \delta_y, a, x, \tilde{\boldsymbol{n}})$$

which weights the joint complete-case contributions by the inverse joint censoring survival. The augmentation term contributes one term per component of $\boldsymbol{N}$, integrated against the corresponding censoring martingale, with conditional expectations Q_0 evaluated on the participant-specific pattern of at-risk indicators at each integration point. Practical estimation would require estimating the joint conditional distribution $\boldsymbol{N} \mid (A, X, \tilde{Y}, \Delta_Y)$. A data augmentation approach similar to the one described herein could in theory be adapted for this purpose. However, this approach could quickly become computationally prohibitive as the number of negative control event types grows, as one would need to expand censored observations into pseudo-observations encompassing all *combinations* of negative control event times greater than the censoring value. Approximating this approach using multiple imputation from conditional distributions could be an attractive and computationally tractable solution to thsi problem. A complete development is deferred to future work.

6.2.5 Leveraging negative control infections to detect and eliminate bias in immune correlates analyses

An important objective of many vaccine studies is identify immune biomarkers that can reliably predict a vaccine's clinical efficacy against disease or infection. Such biomarkers are referred to as immune correlates of protection (CoP).[19–21] Identifying valid CoP is useful for a variety of reasons. Chief among them is a valid CoP can serve as a surrogate endpoint for an infectious disease clinical outcome, which can accelerate vaccine authorization or approval without requiring resource-intensive randomized, controlled phase 3 trials. Moreover, a valid CoP can help elucidate a vaccine's immunological mechanism of action, obtain predictions of VE against new variants, or predict the vaccine's effectiveness in new populations that differ from those typically enrolled in randomized trials.[19–21]

Many statistical frameworks developed for immune correlates analyses rely on the *strong sequential ignorability* assumption, which assumes that the immune biomarker is effectively randomized

within levels of baseline covariates.[21] Since there is no practical way for an experimenter to randomize the immune biomarker level, it is possible that some degree of unmeasured confounding through variables like underlying immune function (U) could distort the relationship. To detect and quantify the magnitude of potential unmeasured confounding on the $S - Y$ relationship, one could use a negative control outcome N that is assumed to be causally unaffected by the vaccine and the immune biomarker but depends on the same sources of unmeasured confounding. Under the NCO assumption as illustrated in Figure 6.2, $S - N$ association can be taken as evidence of unmeasured confounding on the $S - Y$ relationship.

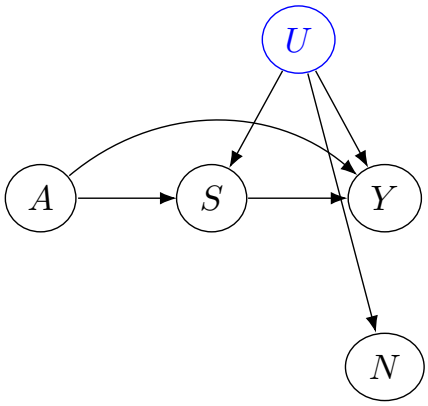

Figure 6.2: Example DAG illustrating how a NCO N could be useful for identifying unmeasured confounding between the association between and immune marker S and the primary infection outcome Y.

A key challenge is that because measuring S is resource-intensive, many vaccine studies will measure S only on a subset of trial participants selected on the basis of available data (namely, the primary outcome), a sampling strategy referred to as two-phase sampling.[22] Common two-phase sampling designs are case-control and case-cohort, which over-sample the breakthrough infections in the vaccine group and randomly sample a subset of the non-cases or a random group of participants. Analyzing the association between an expensive exposure and a secondary outcome when the sampling scheme is informed by the primary outcome is a rich area of research, for which some prior work may be relevant.[23–30]

Bibliography

[1] Michael L Jackson and Jennifer C Nelson. The test-negative design for estimating influenza vaccine effectiveness. *Vaccine*, 31(17):2165–2168, April 2013.

[2] Christopher B Boyer, Kendrick Qijun Li, Xu Shi, and Eric J Tchetgen Tchetgen. Identification and estimation of vaccine effectiveness in the test-negative design under equi-confounding. *Epidemiology*, 37(1):77–87, January 2026.

[3] Ted Westling, Peter Gilbert, and Marco Carone. Causal isotonic regression. *J. R. Stat. Soc. Series B Stat. Methodol.*, 82(3):719–747, July 2020.

[4] Ted Westling and Marco Carone. A unified study of nonparametric inference for monotone functions. *Ann. Stat.*, 48(2):1001–1024, April 2020.

[5] Tara M Babu, Leora R Feldstein, Sharon Saydah, Zachary Acker, Cassandra L Boisvert, Melissa Briggs-Hagen, Marco Carone, Amanda Casto, Sarah N Cox, Brenna Ehmen, Janet A Englund, Stephen P Fortmann, Collrane J Frivold, Holly Groom, Peter D Han, Jennifer L Kuntz, Tina Lockwood, Claire M Midgley, Richard A Mularski, Tara Ogilvie, Sacha L Reich, Mark A Schmidt, Ning Smith, Lea Starita, Jeremy Stone, Meredith Vandermeer, Ana A Weil, Caitlin R Wolf, Helen Y Chu, and Allison L Naleway. CASCADIA: a prospective community-based study protocol for assessing SARS-CoV-2 vaccine effectiveness in children and adults using a remote nasal swab collection and web-based survey design. *BMJ Open*, 13(7):e071446, July 2023.

[6] Leora R Feldstein, Amadea Britton, Lauren Grant, Ryan Wiegand, Jasmine Ruffin, Tara M Babu, Melissa Briggs Hagen, Jefferey L Burgess, Alberto J Caban-Martinez, Helen Y Chu, Katherine D Ellingson, Janet A Englund, Kurt T Hegmann, Zuha Jeddy, Adam S Lauring, Karen Lutrick, Emily T Martin, Clare Mathenge, Jennifer Meece, Claire M Midgley, Arnold S Monto, Gabriella Newes-Adeyi, Leah Odame-Bamfo, Lauren E W Olsho, Andrew L Phillips, Ramona P Rai, Sharon Saydah, Ning Smith, Laura Steinhardt, Harmony Tyner, Meredith Vandermeer, Molly Vaughan, Sarang K Yoon, Manjusha Gaglani, and Allison L Naleway.

Effectiveness of bivalent mRNA COVID-19 vaccines in preventing SARS-CoV-2 infection in children and adolescents aged 5 to 17 years. *JAMA*, 331(5):408–416, February 2024.

[7] Leora R Feldstein, Jasmine Ruffin, Ryan Wiegand, Lauren Grant, Tara M Babu, Melissa Briggs-Hagen, Jefferey L Burgess, Alberto J Caban-Martinez, Helen Y Chu, Katherine D Ellingson, Janet A Englund, Kurt T Hegmann, Zuha Jeddy, Jennifer Kuntz, Adam S Lauring, Karen Lutrick, Emily T Martin, Clare Mathenge, Jennifer Meece, Claire M Midgley, Arnold S Monto, Allison L Naleway, Gabriella Newes-Adeyi, Leah Odame-Bamfo, Lauren E W Olsho, Andrew L Phillips, Ramona P Rai, Sharon Saydah, Ning Smith, Harmony Tyner, Molly Vaughan, Ana A Weil, Sarang K Yoon, Amadea Britton, and Manjusha Gaglani. Protection from COVID-19 vaccination and prior SARS-CoV-2 infection among children aged 6 months-4 years, united states, september 2022-april 2023. *J. Pediatric Infect. Dis. Soc.*, 14(1):piae121, January 2025.

[8] Leora R Feldstein, Jasmine Ruffin, Ryan E Wiegand, Craig B Borkowf, Jade James-Gist, Tara M Babu, Melissa Briggs-Hagen, James Chappell, Helen Y Chu, Janet A Englund, Jennifer L Kuntz, Adam S Lauring, Natalie Lo, Marco Carone, Christina Lockwood, Emily T Martin, Claire M Midgley, Arnold S Monto, Allison L Naleway, Tara Ogilvie, Sharon Saydah, Mark A Schmidt, Jonathan E Schmitz, Ning Smith, Ine Sohn, Lea Starita, H Keipp Talbot, Ana A Weil, and Carlos G Grijalva. Effectiveness of mRNA COVID-19 vaccines and hybrid immunity in preventing SARS-CoV-2 infection and symptomatic COVID-19 among adults in the united states. *J. Infect. Dis.*, 231(4):e743–e753, April 2025.

[9] M. Lipsitch, E. Goldstein, G.T. Ray, and B. Fireman. Depletion-of-susceptibles bias in influenza vaccine waning studies: how to ensure robust results. *Epidemiology & Infection*, 2019.

[10] G Thomas Ray, Ned Lewis, Nicola P Klein, Matthew F Daley, Marc Lipsitch, and Bruce Fireman. Depletion-of-susceptibles bias in analyses of intra-season waning of influenza vaccine effectiveness. *Clin. Infect. Dis.*, 70(7):1484–1486, March 2020.

[11] Jonathan Fintzi and Dean Follmann. Assessing vaccine durability in randomized trials following placebo crossover. *Stat. Med.*, 40(27):5983–6007, November 2021.

[12] Matias Janvin and Mats J Stensrud. Quantification of vaccine waning as a challenge effect. *J. Am. Stat. Assoc.*, pages 1–11, December 2024.

[13] Ahyoung Lim, Freya M Shearer, Kara Sewalk, David M Pigott, Joseph Clarke, Azhar Ghouse, Ciara Judge, Hyolim Kang, Jane P Messina, Moritz U G Kraemer, Katy A M Gaythorpe, William M de Souza, Elaine O Nsoesie, Michael Celone, Nuno Faria, Sadie J Ryan, Ingrid B Rabe, Diana P Rojas, Simon I Hay, John S Brownstein, Nick Golding, and Oliver J Brady. The overlapping global distribution of dengue, chikungunya, zika and yellow fever. *Nat. Commun.*, 16(1):3418, April 2025.

[14] Malaria Atlas Project. Global plasmodium falciparum malaria incidence rate, 2024. `https://data.malariaatlas.org`, 2024. Accessed: 2026-04-23.

[15] Benjamin F Arnold, Francois Rerolle, Christine Tedijanto, Sammy M Njenga, Mahbubur Rahman, Ayse Ercumen, Andrew Mertens, Amy J Pickering, Audrie Lin, Charles D Arnold, Kishor Das, Christine P Stewart, Clair Null, Stephen P Luby, John M Colford, Jr, Alan E Hubbard, and Jade Benjamin-Chung. Geographic pair matching in large-scale cluster randomized trials. *Nat. Commun.*, 15(1):1069, February 2024.

[16] Jennifer P Collins and Andi L Shane. Infections associated with group childcare. In *Principles and Practice of Pediatric Infectious Diseases*, pages 25–32.e3. Elsevier, 2018.

[17] Lawrence Corey, Peter B. Gilbert, Michal Juraska, David C. Montefiori, Lynn Morris, Shelly T. Karuna, Srilatha Edupuganti, Nyaradzo M. Mgodi, Allan C. deCamp, Erika Rudnicki, Yunda Huang, Pedro Gonzales, Robinson Cabello, Catherine Orrell, Javier R. Lama, Fatima Laher, Erica M. Lazarus, Jorge Sanchez, Ian Frank, Juan Hinojosa, Magdalena E. Sobieszczyk, Kyle E. Marshall, Pamela G. Mukwekwerere, Joseph Makhema, Lindsey R. Baden, James I. Mullins, Carolyn Williamson, John Hural, M. Juliana McElrath, Carter Bentley, Simbarashe Takuva, Margarita M. Gomez Lorenzo, David N. Burns, Nicole Espy, April K. Randhawa, Nidhi Kochar, Estelle Piwowar-Manning, Deborah J. Donnell, Nirupama Sista, Philip Andrew, James G. Kublin, Glenda Gray, Julie E. Ledgerwood, John R. Mascola, and Myron S. Cohen. Two Randomized Trials of Neutralizing Antibodies to Prevent HIV-1 Acquisition. *New England*

Journal of Medicine, 384(11):1003–1014, March 2021. ISSN 0028-4793, 1533-4406. doi: 10.1056/NEJMoa2031738. URL http://www.nejm.org/doi/10.1056/NEJMoa2031738.

[18] Srilatha Edupuganti, Nyaradzo Mgodi, Shelly T. Karuna, Philip Andrew, Erika Rudnicki, Nidhi Kochar, Allan deCamp, Robert De La Grecca, Maija Anderson, Carissa Karg, India Tindale, Elizabeth Greene, Gail B. Broder, Jonathan Lucas, John Hural, Jorge A. Gallardo-Cartagena, Pedro Gonzales, Ian Frank, Magdalena Sobieszczyk, Margarita M. Gomez Lorenzo, David Burns, Peter L. Anderson, Maurine D. Miner, Julie Ledgerwood, John R. Mascola, Peter B. Gilbert, Myron S. Cohen, Lawrence Corey, and the HVTN 704/HPTN 085 study group. Feasibility and Successful Enrollment in a Proof-of-Concept HIV Prevention Trial of VRC01, a Broadly Neutralizing HIV-1 Monoclonal Antibody. *JAIDS Journal of Acquired Immune Deficiency Syndromes*, 87(1):671–679, May 2021. ISSN 1525-4135. doi: 10.1097/QAI.0000000000002639. URL https://journals.lww.com/10.1097/QAI.0000000000002639.

[19] Stanley A Plotkin. Vaccines: correlates of vaccine-induced immunity. *Clin. Infect. Dis.*, 47(3):401–409, August 2008.

[20] Stanley A Plotkin. Correlates of protection induced by vaccination. *Clin. Vaccine Immunol.*, 17(7):1055–1065, July 2010.

[21] Peter B Gilbert, Youyi Fong, Nima S Hejazi, Avi Kenny, Ying Huang, Marco Carone, David Benkeser, and Dean Follmann. Four statistical frameworks for assessing an immune correlate of protection (surrogate endpoint) from a randomized, controlled, vaccine efficacy trial. *Vaccine*, 42(9):2181–2190, April 2024.

[22] Peter B Gilbert, Youyi Fong, Avi Kenny, and Marco Carone. A controlled effects approach to assessing immune correlates of protection. *Biostatistics*, 24(4):850–865, October 2023.

[23] A J Lee, L McMurchy, and A J Scott. Re-using data from case-control studies. *Stat. Med.*, 16(12):1377–1389, June 1997.

[24] Marie Reilly, Anna Torrång, and Asa Klint. Re-use of case-control data for analysis of new outcome variables. *Stat. Med.*, 24(24):4009–4019, December 2005.

[25] Yannan Jiang, Alastair J Scott, and Chris J Wild. Secondary analysis of case-control data. *Stat. Med.*, 25(8):1323–1339, April 2006.

[26] David B Richardson, Peter Rzehak, Jochen Klenk, and Stephan K Weiland. Analyses of case-control data for additional outcomes. *Epidemiology*, 18(4):441–445, July 2007.

[27] Genevieve M Monsees, Rulla M Tamimi, and Peter Kraft. Genome-wide association scans for secondary traits using case-control samples. *Genet. Epidemiol.*, 33(8):717–728, December 2009.

[28] Jiawei Wei, Raymond J Carroll, Ursula U Müller, Ingrid Van Keilegom, and Nilanjan Chatterjee. Robust estimation for homoscedastic regression in the secondary analysis of case-control data. *J. R. Stat. Soc. Series B Stat. Methodol.*, 75(1):185–206, January 2013.

[29] Eric J Tchetgen Tchetgen. A general regression framework for a secondary outcome in case-control studies. *Biostatistics*, 15(1):117–128, January 2014.

[30] Yinghao Pan, Jianwen Cai, Matthew P Longnecker, and Haibo Zhou. Secondary outcome analysis for data from an outcome-dependent sampling design. *Stat. Med.*, 37(15):2321–2337, July 2018.

Appendix A

SUPPLEMENTARY MATERIALS (CHAPTER 2)

A.1 Stratification and adjustment for post-baseline variables

It is not common practice to stratify or adjust analyses for variables measured post-baseline due to the possibility of producing selection bias. However, we advocate that selection bias can be avoided if the post-baseline variable is unaffected by treatment. These ideas are pertinent to some Objective 3 analyses.

Frangakis and Rubin [1] illustrated that treatment effects stratified by post-baseline covariates are not necessarily well-defined causal effects. However, causal effects within "principal strata" defined by the joint potential outcome values are well-defined, because membership in a principal stratum is unaffected by treatment. Since potential outcomes under both treatment and control interventions are not typically observed, principal strata causal effects are usually only partially identified. However, we note that when the post-baseline variable is unaffected by treatment, i.e. when N is a negative control outcome satisfying $N(1) = N(0) = N$, that causal effects within principal strata of the NCO are point-identified. Hence, stratification by an NCO represents a well-defined causal estimand not subject to selection bias.

With regards to adjustment for a post-baseline variable, we refer to the work by Rosenbaum [2], who describes how adjustment for a variable measured after treatment provision can result in selection bias if the variable is affected by treatment. Hence, the sufficient and necessary condition allowing unbiased estimation of treatment effects in a randomized experiment is that the analyst only adjusts for variables that are not affected by treatment. In describing causal analyses using the potential outcome framework, Rubin [3] says that adjustment covariates "take their values before the treatment assignment or, more generally, simply cannot be affected by the treatment". Hence, adjustment for a post-baseline variable avoids bias as long as the variable is unaffected by treatment. This includes NCOs as a special case.

A.2 "Validating" candidate negative control endpoints

To "validate" a candidate negative control outcome, we would like to confirm that it is (i) unaffected by the treatment and yet correlated with the primary outcome (presumably through unmeasured common causes)

A.2.1 Evaluating null effect

Our first objective was to test if a candidate NCO was unaffected by mRNA-1273 vaccination using data from the randomized, blinded phase of COVE. Let $A \in \{0, 1\}$ be a binary indicator of randomization to placebo or mRNA-1273 vaccine. Define $N \geq 0$ as a random variable denoting time to a candidate NCO in the blinded period. The time origin will be from the date of enrollment for each participant in accordance with a time-scale in study time. Cases will be counted 14 days after receipt of last dose of mRNA-1273 or placebo. Let $N(a)$ refer to the potential outcomes given hypothetical assignment to treatment $A = a$. Let $P_0(N(a) < t)$ refer to the causal cumulative incidence of the candidate NCO t-days post-baseline.

Our target estimands are the causal VE against the candidate NCO endpoint at time $t_0 = 180$ days, or one minus the ratio of cumulative risks of each outcome between mRNA-1273 and placebo arms.

$$\text{VE}^N(t_0) = 1 - \frac{P_0(N(1) \leq t_0)}{P_0(N(0) < t_0)} \tag{A.1}$$

We will use the method of Westling et al.[4] to estimate the treatment-specific survival functions as implemented in the R package 'CFSurvival'. To estimate the target estimands in a consistent and asymptotically normal fashion, we require consistent estimators of nuisance parameters (the treatment probability, conditional survival functions for each endpoint, and a conditional survival function for the censoring mechanism in each randomization arm). The empirical treatment probability will be used for the first nuisance parameter. The conditional survival functions for the candidate NCO and censoring were estimated using two 10-fold cross-fit SuperLearners with a library including Kaplan-Meier estimators, Cox models, parametric regression models, generalized additive models, and survival random forests and three kinds of variable screening procedures. Inference on

the target estimands can be achieved by combining estimates of the treatment-specific cumulative incidence using the delta method. All estimators of the treatment-arm specific cumulative incidence functions will adjust for the set of baseline covariates described above to relax assumptions on the censoring mechanism and improve precision in estimates of the target estimands.

The effect of mRNA-1273 on each candidate NCO endpoint was evaluated on the basis of point estimates of the target estimands in Equation A.1 and 95% confidence intervals. "No effect" corresponds to a VE not significantly different from zero. A VE that significantly differs from zero suggests invalidity of the NCO.

A.2.2 Evaluating association

The second objective was to examine correlation between the candidate negative control and primary outcome. A Cochran-Mantel-Haenszel odds ratio stratified by randomization arm was conducted to provide a single-number summary of the association between binary indicators of the candidate NCO and COVID-19 infection outcomes in the blinded phase among the per-protocol cohort. A 95% exact confidence interval for the odds ratio will quantify uncertainty in the association.

As a supporting analysis, we estimating cumulative incidence curves for COVID-19 stratified by randomization arm and the candidate NCO infection status during the blinded phase. While it is not typical to stratify an analysis by a post-baseline variable, it is acceptable practice if the variable is unaffected by treatment condition.[3] Let $Y(a)$ refer to the potential time-to-COVID-19 under assignment to treatment $A = a$. Let $t > 0$ be some post-baseline time. Let $\Delta^N := I(N \leq C) \in \{0, 1\}$ refer to an indicator of whether a negative control infection occurred before censoring. For all combinations $a \in \{0, 1\}$ and $\delta_n \in \{0, 1\}$, we estimated

$$P_0(Y(a) \leq t | \Delta^N = \delta^N) \tag{A.2}$$

For $\delta^N \in \{0, 1\}$. As before, we use the method of Westling et al.[4] to estimate the treatment-and-strata-specific cumulative incidence functions for COVID-19 using the 'CFsurvival' package while adjusting for sex, URM status, randomization stratum, and continuous baseline COVID-19 risk score. The treatment propensity will be estimated using the empirical proportion of indi-

viduals randomized to mRNA-1273 in the strata defined by randomization and non-SARS-CoV-2 pathogen-positive ARI status. Nuisance parameters (conditional survival functions for COVID-19 and censoring times) will be estimated using 10-fold cross-fit SuperLearners as described above. We will compute the cumulative incidences stratified by non-SARS-CoV-2 pathogen-positive ARI indicators from $0 \leq t \leq 180$.

A.3 Supplementary Tables and Figures

A.3.1 Tables

Pathogen	**Placebo (n=15,166)**	**mRNA-1273 (n=15,181)**	**Overall (n=30,347)**
SARS-CoV-2-Negative ARI	829	758	1,587
Non-SARS-CoV-2 pathogen-positive ARI	252	261	513
Human Rhinovirus/Enterovirus	248	250	498
Coronavirus NL63	1	4	5
Adenovirus	2	1	3
Bordetella pertussis	0	2	2
Coronavirus HKU1	1	1	2
Coronavirus OC43	0	2	2
Parainfluenza Virus 1	0	1	1

Table A.1: Counts of candidate NCO endpoints in blinded phase among the Full analysis set (FAS) and stratified by randomization arm. Case counts for specific pathogens with at least one infection during blinded phase are shown.

Pathogen	Open-label Phase (Unblinding–September 22, 2021)			Booster Phase (September 23, 2021–April 5, 2022)		Overall
	Crossover mRNA-1273 (n=12,649)	Immediate mRNA-1273 (n=14,662)	Declined mRNA-1273 (n=1,720)	Received Booster (n=19,597)	Declined Booster (n=9,434)	(n=29,031)
Non-SARS-CoV-2 pathogen-positive ARI	402	450	6	834	62	1,754
Human Rhinovirus/Enterovirus	247	268	4	388	36	943
Coronavirus OC43	49	56	1	102	4	212
Respiratory Syncytial Virus	35	39	0	67	10	151
Coronavirus 229E	19	26	0	55	3	103
Parainfluenza Virus 3	27	23	0	22	1	73
Human metapneumovirus	0	3	0	54	5	62
Coronavirus NL63	18	17	1	18	0	54
Influenza A	0	0	0	39	0	39
Influenza A subtype H3	0	0	0	39	0	39
Parainfluenza Virus 2	6	9	0	17	0	32
Parainfluenza Virus 4	0	3	0	14	1	18
Adenovirus	1	4	0	9	0	14
Coronavirus HKU1	0	2	0	8	2	12
Bordetella pertussis	0	0	0	1	0	1
Parainfluenza Virus 1	0	0	0	1	0	1

Table A.2: Counts of candidate NCO endpoints in blinded phase among the Full analysis set (FAS) and stratified by randomization arm. Case counts for specific pathogens with at least one infection during blinded phase are shown.

Variable	Level	Incidence Rate (95% CI)	Incidence Rate Ratio (95% CI)
Randomization Arm	Placebo	4.29 (3.77, 4.85)	1 (reference)
	mRNA	4.33 (3.82, 4.88)	1.01 (0.849, 1.20)
Randomization Stratum	Age [18,65) and Not at Risk	5.18 (4.66, 5.75)	1 (reference)
	Age [18,65) and at Risk	4.44 (3.57, 5.45)	0.856 (0.68, 1.47)
	≥65 Years	2.15 (1.65, 2.75)	0.414 (0.317, 0.541)
Sex	Female	4.99 (4.43, 5.61)	1 (reference)
	Male	3.69 (3.22, 4.19)	0.739 (0.621, 0.879)
Racial/Ethnic URM	Not URM	4.15 (3.56, 4.81)	1 (reference)
	URM	4.39 (3.94, 4.89)	1.06 (0.882, 1.27)

Table A.3: Incidence rates per 100 PYFU, incidence rate ratios, and 95% Poisson exact confidence intervals (CIs) for non-SARS-CoV-2 pathogen-positive ARI in COVE's blinded phase among FAS (N=30,347).

Variable	Level	Incidence Rate (95% CI)	Incidence Rate Ratio (95% CI)
Randomization Strata	Age [18,65) and Not at Risk	6.76 (6.37, 7.17)	1 (reference)
	Age [18,65) and at Risk	5.76 (5.11, 6.47)	0.853 (0.748, 0.971)
	≥65 Years	4.65 (4.18, 5.16)	0.688 (0.61, 0.775)
Sex	Female	6.40 (5.99, 6.84)	1 (reference)
	Male	5.69 (5.32, 6.09)	0.889 (0.81, 0.977)
Racial/Ethnic URM	Not URM	5.87 (5.52, 6.23)	1 (reference)
	URM	6.31 (5.85, 6.81)	1.08 (0.978, 1.19)

Table A.4: Incidence rates per 100 PYFU, incidence rate ratios, and 95% Poisson exact confidence intervals (CIs) for non-SARS-CoV-2 pathogen-positive ARI in COVE's open-label and booster phases among FAS participants who began the open-label phase (N=29,031).

A.3.2 Figures

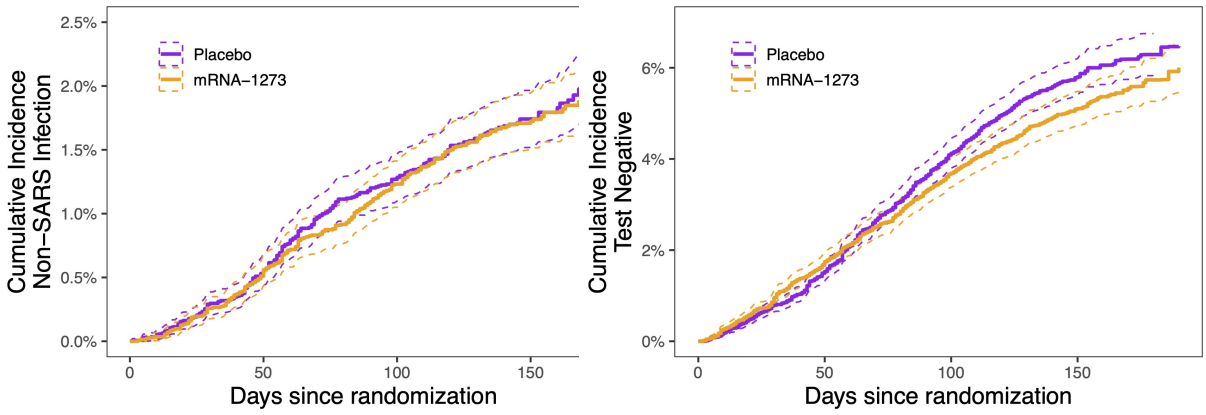

Figure A.1: Cumulative incidence of non-SARS-CoV-2 pathogen-positive ARI (left) and SARS-CoV-2 negative ARI (right) counted after randomization date. Curves were estimated from the blinded phase per- protocol population. Cumulative incidence estimates adjusted for sex, racial/ethnic URM status, randomization strata, and baseline risk score. 95% pointwise confidence intervals shown.

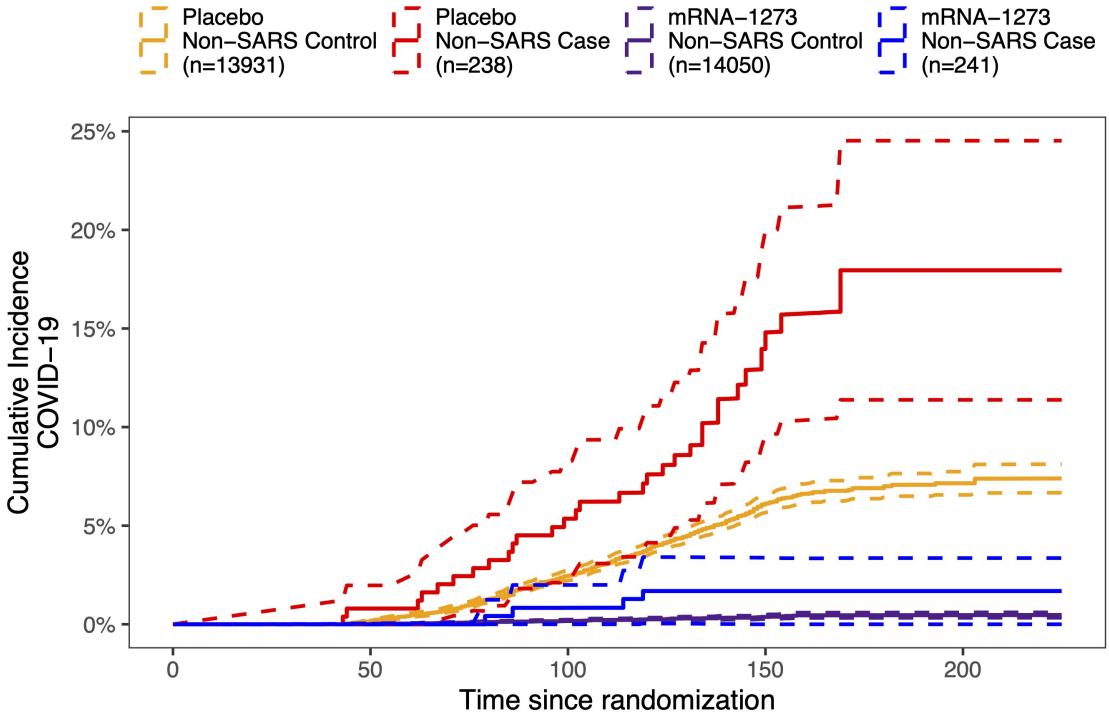

Figure A.2: Cumulative incidence of COVID-19 in the blinded phase per-protocol cohort stratified by randomization arm and indicator of non-SARS-CoV-2 pathogen-positive ARI during the blinded phase. "Non-SARS cases" refer to participants who acquired a non-SARS-CoV-2 respiratory infection during COVE blinded phase; remaining participants were designated "Non-SARS controls". Curves are adjusted for sex, randomization strata, racial/ethnic URM status, and baseline risk score. 95% pointwise Wald CIs shown.

Bibliography

[1] CE Frangakis and DB Rubin. Principal stratification in causal inference. *Biometrics*, 58:21–29, 2002.

[2] PR Rosenbaum. The consequence of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society, Series A*, 147:656–666, 1984.

[3] DB Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100:322–331, 2005.

[4] Ted Westling, Alex Luedtke, Peter B. Gilbert, and Marco Carone. Inference for Treatment-

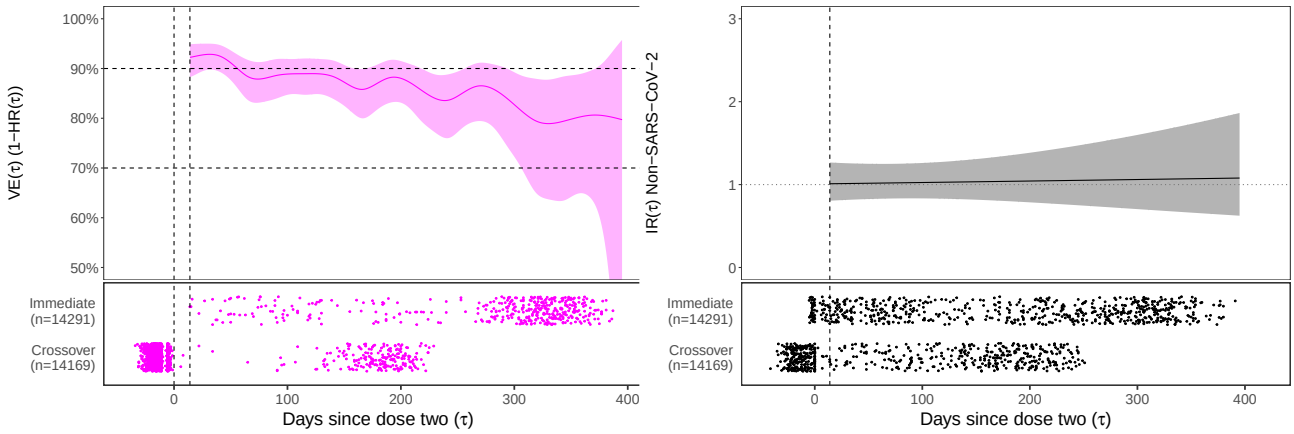

Figure A.3: (Left) Time-varying Vaccine efficacy (VE) of the two-dose mRNA-1273 vaccination series against COVID-19 during the blinded and open-label phases (after placebo crossover). VE was defined as one minus the hazard ratio which varied in time-since-vaccination according to a cubic P-spline. VE was estimated starting after a blackout period between first and 14-days after second dose to capture full effect of immunization. (Right) Time-varying intensity ratio of non-SARS-CoV-2 pathogen-positive ARI outcomes during the blinded and open-label phases (after placebo crossover). Varying IR was modeled using cubic P-spline. In all analyses, different baseline hazards were used for different combinations of sex, racial/ethnic URM status, blind status, and randomization strata. Estimates adjusted for continuous baseline COVID-19 risk score. 95% pointwise Wald CIs shown.

Specific Survival Curves Using Machine Learning. *Journal of the American Statistical Association*, 0(0):1–13, 2023. ISSN 0162-1459. doi: 10.1080/01621459.2023.2205060. URL `https://doi.org/10.1080/01621459.2023.2205060`.

Appendix B

SUPPLEMENTARY MATERIALS (CHAPTER 3)

B.1 Identification Results

B.1.1 Identification of time-varying, frailty-standardized VE parameter

Let $F^1_{U(t)}$ refer to the distribution of frailties among vaccinated survivors: i.e., those with $\{T \geq t, V = v\}$. Recall our statistical parameter is the time-varying VE standardized to the frailty distribution in this subgroup.

$$\text{VE}_1(t,v) := 1 - \frac{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V = v, U(t))]}{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V > t, U(t))]}$$

Which avoids depletion of susceptibles bias and confounding by standardizing the numerator and denominator to a common frailty distribution. The critical challenge is that the frailties are unobserved. Hence, the core aspect of our identification task centers around finding an way to indirectly identify the frailty distributions and their impact on different subgroups of participants.

As a starting point, we write out the observed marginal hazard ratio of vaccine-preventable infection.

$$\begin{aligned}
\text{HR}_1(t,v) &:= \frac{h_1(t \mid V = v)}{h_1(t \mid V > t)} \\
&= \frac{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V = v, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_1(t \mid V > t, U(t))]} \\
&= \underbrace{\frac{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V = v, U(t))]}{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V > t, U(t))]}}_{\equiv 1 - \text{VE}_1(t,v)} \times \frac{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V > t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_1(t \mid V > t, U(t))]}
\end{aligned}$$

Where the first equality is definitional, the second follows from Tower law, and the third from multiplying and dividing by a common term. Hence, the observed hazard ratio of vaccine preventable infection equals our target parameter times a bias term caused by marginalizing over two frailty distributions that may differ due to selection bias (i.e., depletion of susceptibles bias) or confounding

in the case of an observational study.

We can apply very similar arguments to the observed time-varying hazard ratio for vaccine-irrelevant infections.

$$\begin{aligned}\mathrm{HR}_0(t,v) &:= \frac{h_0(t \mid V=v)}{h_0(t \mid V>t)} \\ &= \frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V=v, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_0(t \mid V>t, U(t))]} \\ &= \frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V>t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_0(t \mid V>t, U(t))]}\end{aligned}$$

Where the first equality is definitional, the second follows from tower law, and the third follows from the NCO assumption (Assumption 2), which claims that the frailty-conditional hazard of vaccine-irrelevant infections is independent of vaccination date. Now the vaccine-irrelevant hazard ratio equals a term that may quantify the imbalance in frailty distributions between vaccinated and unvaccinated survivors that appears in the hazard ratio of vaccine-preventable pathogens.

Assumption 4 presumes that the effect of the frailty imbalance on the hazard ratio of vaccine-preventable infections can be quantified by the frailty imbalance on the hazard ratio of vaccine-irrelevant infections.

$$\frac{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V>t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_1(t \mid V>t, U(t))]} \overset{\text{A4}}{=} \frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V>t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_0(t \mid V>t, U(t))]}$$

This is akin to a hazard ratio equiconfounding assumption commonly used in NCO-based identification strategies.

Under this assumption, the target parameter, $\mathrm{VE}_1(t,v)$ is identified by one minus the ratio of hazard ratios

$$\begin{aligned}1 - \frac{\mathrm{HR}_1(t,v)}{\mathrm{HR}_0(t,v)} &= \mathrm{VE}_1(t,v) \times \frac{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V>t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_1(t \mid V>t, U(t))]} \times \left(\frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V>t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_0(t \mid V>t, U(t))]}\right)^{-1} \\ &\overset{\text{A4}}{=} \mathrm{VE}_1(t,v)\end{aligned}$$

As constructed, $\mathrm{VE}_1(t,v)$ varies in two calendar time inputs, which complicates interpretability.

Assumption 3 stipulates that VE does not depend on calendar time (t, v), but instead depends on a single, interpretable axis: time-since-vaccination $(t - v)$.

$$\mathrm{VE}_1(t-v) \overset{\text{A3}}{=} 1 - \frac{\mathrm{HR}_1(t,v)}{\mathrm{HR}_0(t,v)}$$

So we conclude that the time-varying, frailty-standardized VE is identified by one minus the ratio of hazard ratios.

To complete our identification result, we rewrite the ratio of hazard ratios using the definition of the marginal hazard.

$$\frac{\mathrm{HR}_1(t,v)}{\mathrm{HR}_0(t,v)} = \lim_{\epsilon \downarrow 0} \frac{P(J=1, T \in [t, t+\epsilon) \mid V = v, T \geq t)/P(J=1, T \in [t, t+\epsilon) \mid V > t, T \geq t)}{P(J=0, T \in [t, t+\epsilon) \mid V = v, T \geq t)/P(J=0, T \in [t, t+\epsilon) \mid V > t, T \geq t)}$$

To proceed, can write the probability of failure of type $J = j$ conditional on vaccination date $V = v$ and some infection at time t in terms of the hazard function using Bayes rule.

$$\begin{aligned} P(J = j \mid T = t, V = v) &= \lim_{\epsilon \downarrow 0} P(J = j \mid T \in [t, t+\epsilon), V = v) \\ &= \lim_{\epsilon \downarrow 0} \frac{P(J = j, T \in [t, t+\epsilon) \mid V = v, T \geq t)}{P(T \in [t, t+\epsilon) \mid V = v, T \geq t)} \end{aligned}$$

Hence, multiplying and diving the ratio of hazard ratios by $P(T \in [t, t+\epsilon) \mid V = v, T \geq t)$ and $P(T \in [t, t+\epsilon) \mid V > t, T \geq t)$ yields

$$\frac{\mathrm{HR}_1(t,v)}{\mathrm{HR}_0(t,v)} = \frac{P(J=1 \mid T = t, V = v)/P(J=0 \mid T = t, V = v)}{P(J=1 \mid T = t, V > t)/P(J=0 \mid T = t, V > t)}$$

Which is the conditional odds ratio of infection causes among those infected on date t between the vaccinated and unvaccinated group. Hence, the time-varying, frailty-standardized VE is equal to one minus the odds ratio.

$$\mathrm{VE}_1(t-v) = 1 - \frac{P(J=1 \mid T = t, V = v)/P(J=0 \mid T = t, V = v)}{P(J=1 \mid T = t, V > t)/P(J=0 \mid T = t, V > t)}$$

B.1.2 Sensitivity analysis

We can parametrize violations of Assumption 2 (NCO) by allowing a user-specified hazard ratio

$$\frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V = v, U(t))]}{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V > t, U(t))]} = \xi(t - v) \neq 1$$

Under this assumption, the marginal time-varying hazard ratio of irrelevant infections is given by

$$\begin{aligned}
\mathrm{HR}_0(t, v) &= \frac{h_0(t \mid V = v)}{h_0(t \mid V > t)} \\
&= \frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V = v, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_0(t \mid V > t, U(t))]} \\
&= \frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V = v, U(t))]}{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V > t, U(t))]} \times \frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V > t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_0(t \mid V > t, U(t))]} \\
&= \xi(t - v) \times \frac{E_{U(t)\sim F^1_{U(t)}}[h_0(t \mid V > t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_0(t \mid V > t, U(t))]}
\end{aligned}$$

The time-varying hazard ratio of vaccine-preventable infections is writeable as.

$$\begin{aligned}
\mathrm{HR}_1(t, v) &:= \frac{h_1(t \mid V = v)}{h_1(t \mid V > t)} \\
&= (1 - \mathrm{VE}_1(t, v)) \times \frac{E_{U(t)\sim F^1_{U(t)}}[h_1(t \mid V > t, U(t))]}{E_{U(t)\sim F^0_{U(t)}}[h_1(t \mid V > t, U(t))]}
\end{aligned}$$

Hence,

$$\frac{\mathrm{HR}_1(t, v)}{\mathrm{HR}_0(t, v)} = (1 - \mathrm{VE}_1(t, v)) \times [\xi(t - v)]^{-1}.$$

Some algebraic manipulation shows that the target parameter is identified up to the analyst-defined constant, $\xi(t - v)$, which attenuates the estimate of the vaccine effectiveness.

$$\mathrm{VE}_1(t, v) = 1 - \left(\xi(t - v) \times \frac{P(J = 1 \mid T = t, V = v)/P(J = 0 \mid T = t, V = v)}{P(J = 1 \mid T = t, V = \infty)/P(J = 0 \mid T = t, V = \infty)}\right)$$

We can parametrize violations of bias equivalence (Assumption 4) by assuming equivalence up

to a multiplicative factor $\beta(t)$.

$$\frac{E_{U(t)\sim F^{(1)}_{U(t)}}[h_1(t \mid V > t, U(t))]}{E_{U(t)\sim F^{(0)}_{U(t)}}[h_1(t \mid V > t, U(t))]} = \beta(t) \times \frac{E_{U(t)\sim F^{(1)}_{U(t)}}[h_0(t \mid V > t, U(t))]}{E_{U(t)\sim F^{(0)}_{U(t)}}[h_0(t \mid V > t, U(t))]}$$

Then, the time-varying vaccine effectiveness is identified up to this constant by the ratio of hazard ratios.

$$\frac{\mathrm{HR}_1(t, v)}{\mathrm{HR}_0(t, v)} = (1 - \mathrm{VE}_1(t, v)) \cdot \beta(t)$$

Implying that

$$\mathrm{VE}_1(t - v) = 1 - \left(\beta(t)^{-1} \times \frac{P(J = 1 \mid T = t, V = v)/P(J = 0 \mid T = t, V = v)}{P(J = 1 \mid T = t, V = \infty)/P(J = 0 \mid T = t, V = \infty)}\right)$$

B.2 Equivalence between general identification result and partial/profile likelihood under multiplicative frailty model

Recall the following multiplicative hazard frailty models which we used to motivate our approach.

$$h_1(t) = h_{01}(t)U(t)\exp\{I(v \le t)f(t - v)\}$$
$$h_0(t) = h_{00}(t)U(t)$$

Below, we will show that partial/profile likelihood techniques suitable for inference in the above multiplicative frailty hazard models identically reproduce our general identification result for the time-varying, frailty-standardized VE

$$\begin{aligned}\mathrm{VE}_1(t - v) &:= 1 - \frac{E_{U(t)\in F^1_{U(t)}}[h_1(t \mid V = v, U(t))]}{E_{U(t)\in F^1_{U(t)}}[h_1(t \mid V > t, U(t))]} \\ &= 1 - \frac{P_0(J = 1 \mid T = t, V = v)/P(J = 0 \mid T = t, V = v)}{P_0(J = 1 \mid T = t, V > t)/P(J = 0 \mid T = t, V > t)}.\end{aligned}$$

To proceed, we note that we can convert the hazards of vaccine-preventable and irrelevant hazards to the same calendar time scale via a change-of-scale parameter $\alpha_0(t) = \log\{h_{01}(t)/h_{00}(t)\}$.

$$h_1(t) = h_{00}(t)U(t)\exp\{I(v \le t)f(t - v) + \alpha_0(t)\}$$

$$h_0(t) = h_{00}(t)U(t)$$

Our goal is to develop a way to identify and estimate $f(\cdot)$ which encodes how VE varies in time-since-vaccination. However, a key challenge is how to avoid specifying the participant-level baseline hazard function: $h_{0i}(t) := h_{00}(t)U_i(t)$. Since the number of baseline hazards (i.e., nuisance parameters) grows linearly with the sample size, attempts to direct model the baseline hazards are infeasible. To identify and estimate $f(\cdot)$ while eliminating the need to model the (participant-level) baseline hazards, we will pursue two classic techniques for nuisance parameter elimination – a partial likelihood and a profile likelihood approach. In accordance with a classic result, we can show that both the profile and partial likelihood approaches are equivalent to one another. Interestingly, we can show that these classic nuisance parameter elimination techniques reproduce the main identification result of the time-varying, frailty-standardized VE shown above.

B.2.1 Individually-stratified partial likelihood

To adapt partial likelihood to the setting where each participant has their own baseline hazard, we appeal to partial likelihood techniques developed for survival analysis in twin/match-pair experiments.[1] In such studies, we measure the same survival outcome within pairs of individuals. The key idea is that stratification at the pair-level controls for unmeasured pair-level effects that would otherwise generate bias in naïve, unstratified Cox regression.

To extend this idea to our framework, we recognize that the vaccine-preventable and irrelevant infection outcomes can be conceptualized as "twin" infection outcomes measured within the same participant. Conceptually, we presume that vaccine-preventable and irrelevant infection outcomes measured within a single participant are excellent comparators for one another because they depend on very similar sets of unmeasured risk factors and behaviors. We can borrow the adapt the pair stratification technique developed for twin/matched-pair experiments directly to an "individual/participant-level stratification" when paired infection outcomes are measured within individual patients.

Let T_j refer to the time of infection with outcome type J ($J = 0$ denotes vaccine-irrelevant and $J = 1$ denote vaccine-preventable). Let T refer to the minimum of both infection times. As shorthand, let $Z(t) := I(v \leq t)$ denote current vaccination status. Define the individualized risk set

as $R_i(t) := \{j : T_j \geq t\}$, which captures the outcomes that participant i is at-risk of experiencing at time t. Suppose that we observe n independent and identically distributed (iid) observations $\{(T_i, J_i, V_i)\}_{i=1}^n$. The individually-stratified partial-likelihood under the multiplicative frailty model equals.

$$
\begin{aligned}
L^{\text{partial}}(f, \alpha) &= \prod_{i=1}^{n} \left(\frac{h_{i1}(T_i)}{\sum_{j \in R_i(T_i)} h_{ij}(T_i)} \right)^{J_i} \left(\frac{h_{i0}(T_i)}{\sum_{j \in R_i(T_i)} h_{ij}(T_i)} \right)^{1-J_i} \\
&= \prod_{i=1}^{n} \left(\frac{h_{0i}(T_i) \exp\{Z_i(T_i) f(T_i - V_i) + \alpha(T_i)\}}{h_{0i}(T_i) + h_{0i}(T_i) \exp\{Z_i(T_i) f(T_i - V_i) + \alpha(T_i)\}} \right)^{J_i} \left(\frac{h_{0i}(T_i)}{h_{0i}(T_i) + h_{0i}(T_i) \exp\{Z_i(T_i) f(T_i - V_i) + \alpha(T_i)\}} \right)^{1-J_i} \\
&= \prod_{i=1}^{n} \left(\frac{\exp\{Z_i(T_i) f(T_i - V_i) + \alpha(T_i)\}}{1 + \exp\{Z_i(T_i) f(T_i - V_i) + \alpha(T_i)\}} \right)^{J_i} \left(\frac{1}{1 + \exp\{Z_i(T_i) f(T_i - V_i) + \alpha(T_i)\}} \right)^{1-J_i}
\end{aligned}
$$

Where the first step follows from the definition of the Cox partial likelihood that conditions on the individual participant i, the second follows from substituting the forms of the hazard from the multiplicative frailty models, and the third equality follows from canceling out the participant-level nuisance parameters $h_{0i}(T_i)$ from the numerator and denominator of the likelihood. Recall that $h_{0i}(t) := h_{00}(t)U_i(t)$ – eliminating $h_{0i}(t)$ from the partial likelihood also eliminated the frailty which summarized the potential sources of bias in a naïve, unstratified Cox regression.

Note that the individually-stratified partial likelihood L^{partial} takes the form of a likelihood associated with a conditional logistic regression model. We can read off the probabilities of failing from the vaccine-preventable and irrelevant causes as follows.

$$
\begin{aligned}
P(J = 1 \mid T = t, V = v) &= \frac{\exp\{I(v \leq t) f(t - v) + \alpha_0(t)\}}{1 + \exp\{I(v \leq t) f(t - v) + \alpha_0(t)\}} \\
P(J = 0 \mid T = t, V = v) &= \frac{1}{1 + \exp\{I(v \leq t) f(t - v) + \alpha_0(t)\}}
\end{aligned}
$$

Using some algebra, we can show that the (log) odds ratio of infection causes conditional on infection on date t identifies $f(\cdot)$, the parameter of interest from the multiplicative frailty model.

$$
\log \left(\frac{P(J = 1 \mid T = t, V = v) / P(J = 0 \mid T = t, V = v)}{P(J = 1 \mid T = t, V > t) / P(J = 0 \mid T = t, V > t)} \right) = I(v \leq t) f(t - v)
$$

However, recall that the conditional odds ratio of infection causes identified $\text{VE}_1(t - v)$ in our more general framework. Hence, we have shown that partial likelihood, a classic technique for nuisance parameter elimination in Cox models, reproduces our more general identification result

that did not assume a particular hazard parametrization.

B.2.2 Profile likelihood

A well-known result is that the Cox partial likelihood can also be derived as a profile likelihood.[2] We illustrate the equivalence between partial and profile likelihood approaches, which was used for a similar self-controlled case series approach applied to frailty models in EHR data.[3]

Let T_1 and T_0 refer to the time-to-vaccine-preventable and time-to-vaccine-irrelevant infection respectively. Recall that $h_{0i}(t) = h_{00}(t)U_i(t)$ is the participant-level baseline hazard on the scale of irrelevant infections. Using the hazard frailty models, we can write the likelihood of the full survival data among n individuals in the study as follows.[4]

$$L^{\text{full}}(f,\alpha) = \prod_{i=1}^{n} [h_{0i}(T_{i0})]^{1-J_i} \exp\left(-\int_0^{T_{i0}} h_{0i}(u)du\right) \\ [h_{0i}(T_{i1}) \exp\{Z_i(T_{i1})[f(T_{i1}-V_i)] + \alpha(T_{i1})\}]^{J_i} \\ \exp\left(-\int_0^{T_{i1}} h_{0i}(u) \exp\{Z_i(u)[f(u-V_i)] + \alpha(u)\}du\right)$$

To eliminate the unknown nuisance parameters, we adopt a profile likelihood approach.[2] To begin, we replace the participant-level nuisance function with its finite-dimensional values at the observed event times. This essentially treats the unknown participant-level hazard function as a set of scalar nuisance parameters. We can write the components of the hazard as follows[3,5].

$$h_{0i}(T_{i0}) = (1-J_i)h_{00i}, \qquad h_{0i}(T_{i1}) = J_i h_{01i}, \qquad \int_0^{T_{i0}} h_{0i}(u)du = (1-J_i)h_{0i} + S_i J_i h_{01i}$$

$$\int_0^{T_{i1}} h_{0i}(u) \exp\{Z_i(u)[f(u-V_i,\beta)] + \alpha(u)\}du = J_i h_{01i} \exp\{Z_i(T_{i1})f(T_{i1}-V_i,\beta) + \alpha(T_{i1})\}$$

$$+ P_i(1-J_i)h_{00i} \exp\{Z_i(T_{i0})f(T_{i0}-V_i,\beta) + \alpha(T_{i0})\}$$

Where $P_i := I(T_{i0} \leq T_{i1})$ and $S_i := I(T_{i1} \leq T_{i0})$ are the paired event time ranks. Rewriting the log-likelihood in terms of the new scalar nuisance parameters, h_{00i}, h_{01i}, we obtain a log-likelihood

that can b solved for the nuisances

$$\begin{aligned}\ell^{\text{full}} &= \prod_{i=1}^{n}(1-J_i)\log\{(1-J_i)h_{00i}\} - [(1-J_i)h_{00i} + S_iJ_ih_{01i}] \\ &+ J_i\left(\log\{J_ih_{01i}\} + Z_i(T_{i1})f(T_{i1}-V_i) + \alpha(T_{i1})\right) \\ &- [J_ih_{01i}\exp\{Z_i(T_{i1})f(T_{i1}-V_i)+\alpha(T_{i1})\}] - [P_i(1-J_i)h_{00i}\exp\{Z_i(T_{i0})f(T_{i0}-V_i)+\alpha(T_{i0})\}]\end{aligned}$$

We can construct a score equation with respect to the unknown scalar nuisance parameters, h_{00i}, h_{01i}, and solve for each nuisance parameter with respect to the other variables.

$$\begin{aligned}0 = \frac{\partial \ell_{\text{full}}}{\partial h_{00i}} &= \frac{(1-J_i)}{h_{00i}} - (1-J_i) - P_i(1-J_i)\exp\{Z_i(T_{i0})f(T_{i0}-V_i)+\alpha(T_{i0})\} \\ &\implies h_{00i} = \frac{1}{1+P_i\exp\{Z_i(T_{i0})f(T_{i0}-V_i)+\alpha(T_{i0})\}} \\ 0 = \frac{\partial \ell_{\text{full}}}{\partial h_{01i}} &= -S_iJ_i + \frac{J_i}{h_{01i}} - J_i\exp\{Z_i(T_{i1})f(T_{i1}-V_i)+\alpha(T_{i1})\} \\ &\implies h_{01i} = \frac{1}{S_i+\exp\{Z_i(T_{i1})f(T_{i1}-V_i)+\alpha(T_{i1})\}}\end{aligned}$$

Now we compute the MLEs of the nuisance parameters specified above.

$$h_{0i}(T_{i0}) = \frac{(1-J_i)}{1+P_i\exp\{Z_i(T_{i0})f(T_{i0}-V_i)+\alpha(T_{i0})\}}, \qquad h_{0i}(T_{i1}) = \frac{J_i}{S_i+\exp\{Z_i(T_{i1})f(T_{i1}-V_i)+\alpha(T_{i1})\}}$$

$$\int_0^{T_{i0}} h_{0i}(u)du = \frac{(1-J_i)}{1+P_i\exp\{Z_i(T_{i0})f(T_{i0}-V_i)+\alpha(T_{i0})\}} + \frac{S_iJ_i}{S_i+\exp\{Z_i(T_{i1})f(T_{i1}-V_i)+\alpha(T_{i1})\}}$$

$$\int_0^{T_{i1}} h_{0i}(u)\exp\{Z_i(u)[f(u-V_i)]+\alpha(u)\}du = \frac{J_i\exp\{Z_i(T_{i1})[f(T_{i1}-V_i)]+\alpha(T_{i1})\}}{S_i+\exp\{Z_i(T_{i1})f(T_{i1}-V_i)+\alpha(T_{i1})\}} + \frac{P_i(1-J_i)\exp\{Z_i(T_{i0})[f(T_{i0}-V_i)]+\alpha(T_{i0})\}}{1+P_i\exp\{Z_i(T_{i0})f(T_{i0}-V_i)+\alpha(T_{i0})\}}$$

Combining the cumulative hazards in the manner specified in the full likelihood leads to cancellation

$$\begin{aligned}&\int_0^{T_{i0}} h_{0i}(u)du + \int_0^{T_{i1}} h_{0i}(u)\exp\{Z_i(u)[f(u-V_i)]+\alpha(u)\}du = [1+P_i\exp\{Z_i(T_{i0})f(T_{i0}-V_i)+\alpha(T_{i0})\}] \\ &\times\left(\frac{(1-J_i)}{1+P_i\exp\{Z_i(T_{i0})f(T_{i0}-V_i)+\alpha(T_{i0})\}}\right) + \left(\frac{J_i\times[S+\exp\{Z_i(T_{i1})f(T_{i1}-V_i)+\alpha(T_{i1})\}]}{S_i+\exp\{Z_i(T_{i1})f(T_{i1}-V_i)+\alpha(T_{i1})\}}\right) \\ &= (1-J_i)+J_i\end{aligned}$$

Which does not depend on any parameters of interest. Hence, we can ignore the contributions of the cumulative hazards in computing the likelihood.

Now, we plug the MLEs for the nuisance parameters back into the original full data likelihood to profile out the nuisances. We obtain

$$L^{\text{profile}}(f,\alpha) = \prod_{i=1}^{n} \left(\frac{1}{1 + P_i \exp\{Z_i(T_{i0}) f(T_{i0} - V_i) + \alpha(T_{i0})\}} \right)^{1-J_i} \left(\frac{\exp\{Z_i(T_{i1}) f(T_{i1} - V_i) + \alpha(T_{i1})\}}{S_i + \exp\{Z_i(T_{i1}) f(T_{i1} - V_i) + \alpha(T_{i1})\}} \right)^{J_i}$$

By inspection, we note that the likelihood only depends on the first event time by virtue of the indicators P_i and S_i. Hence, define $T_i = \min(T_{i0}, T_{i1})$. Then we can rewrite the profile likelihood equivalently as

$$L^{\text{profile}}(f,\alpha) = \prod_{i=1}^{n} \left(\frac{1}{1 + \exp\{Z_i(T_i) f(T_i - V_i) + \alpha(T_i)\}} \right)^{1-J_i} \left(\frac{\exp\{Z_i(T_i) f(T_i - V_i) + \alpha(T_i)\}}{1 + \exp\{Z_i(T_i) f(T_i - V_i) + \alpha(T_i)\}} \right)^{J_i}$$

Which is identical to the individually-stratified partial likelihood and therefore equivalent to our general identification result for the time-varying, frailty-standardized VE.

B.3 Connection to Causal VE parameter

B.3.1 Main identification result

Under Assumptions 5-7, we can write the causal quantity $\theta_C(t-v)$ as follows.

$$\begin{aligned}
\theta_C(t-v) &= \lim_{\epsilon \downarrow 0} \frac{P(T_1(v,0) \in [t, t+\epsilon) | T(v,0) \geq t, V = v)}{P(T_1(\infty,t) \in [t, t+\epsilon) | T(v,0) \geq t, V = v)} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(T_1(v,0) \in [t, t+\epsilon) | T(v,0) \geq t, V = v)}{\alpha_0(t) \times P(T_0(\infty,t) \in [t, t+\epsilon) | T(v,0) \geq t, V = v)} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(T_1(v,0) \in [t, t+\epsilon) | T(v,0) \geq t, V = v)}{\alpha_0(t) \times P(T_0(v,t) \in [t, t+\epsilon) | T(v,0) \geq t, V = v)} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(T_1(v,0) \in [t, t+\epsilon) | T(v,0) \geq t, V = v)}{P(T_0(v,0) \in [t, t+\epsilon) | T(v,0) \geq t, V = v)} \times [\alpha_0(t)]^{-1} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(T_1 \in [t, t+\epsilon) | T \geq t, V = v)}{P(T_0 \in [t, t+\epsilon) | T \geq t, V = v)} \times [\alpha_0(t)]^{-1} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(J = 1, T \in [t, t+\epsilon) | T \geq t, V = v)}{P(J = 0, T \in [t, t+\epsilon) | T \geq t, V = v)} \times [\alpha_0(t)]^{-1}
\end{aligned}$$

Where the first equality is restating the definition of $\theta_C(t-v)$, the second equality follows from Assumption 5 (bias equivalence), the third equality from Assumption 6 (NCO assumption), the fourth from applying Assumption 7 (exposure deferral irrelevance) to the denominator, the fifth from consistency of potential outcomes, and the last from rewriting the events $\{T_j \in [t, t+\epsilon)\}$ in

an equivalent way as $\{J = j, T \in [t, t+\epsilon)\}$.

To finish the proof, all that is left is to identify $\alpha_0(t)$.

$$\begin{aligned}
\alpha_0(t) &= \lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty, t) \in [t, t+\epsilon) \mid T(\infty, 0) \geq t, V = \infty)}{P(T_0(\infty, t) \in [t, t+\epsilon) \mid T(\infty, 0) \geq t, V = \infty)} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty, 0) \in [t, t+\epsilon) \mid T(\infty, 0) \geq t, V = \infty)}{P(T_0(\infty, 0) \in [t, t+\epsilon) \mid T(\infty, 0) \geq t, V = \infty)} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(T_1 \in [t, t+\epsilon) \mid T \geq t, V = \infty)}{P(T_0 \in [t, t+\epsilon) \mid T \geq t, V = \infty)} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(J = 1, T \in [t, t+\epsilon) \mid T \geq t, V = \infty)}{P(J = 0, T \in [t, t+\epsilon) \mid T \geq t, V = \infty)}
\end{aligned}$$

Where the first equality follows from Assumption 5 (bias equivalence), the second follows from applying Assumption 7 (exposure deferral irrelevance) to the numerator and denominator, the third follows from applying consistency of potential outcomes, and the last follows by rewriting the events $\{T_j \in [t, t+\epsilon)\}$ in an equivalent way as $\{J = j, T \in [t, t+\epsilon)\}$.

Hence, after cancellation, $\theta_C(t - v)$ is identified by the observable ratio of cause-specific hazard ratios.

$$\theta_C(t - v) = \lim_{\epsilon \downarrow 0} \frac{P(J = 1, T \in [t, t+\epsilon) | T \geq t, V = v)}{P(J = 0, T \in [t, t+\epsilon) | T \geq t, V = v)} \times \left(\frac{P(J = 1, T \in [t, t+\epsilon) \mid T \geq t, V = \infty)}{P(J = 0, T \in [t, t+\epsilon) \mid T \geq t, V = \infty)} \right)^{-1}$$

By applying identical arguments as the previous subsection, we can show that the causal parameter is identified as the odds ratio of infection causes at time $T = t$ between vaccinated and unvaccinated participants.

$$\theta_C(t - v) = \frac{P(J = 1 \mid T = t, V = v)}{P(J = 0 \mid T = t, V = v)} \times \left(\frac{P(J = 1 \mid T = t, V = \infty)}{P(J = 0 \mid T = t, V = \infty)} \right)^{-1}$$

Moreover, suppose we make an additional assumption linking the relative event rates in the "never vaccinated" group to "not yet vaccinated" groups. For all $s > t$, we assume

$$\lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty, 0) \in [t, t+\epsilon) \mid T(\infty, 0) \geq t, V = \infty)}{P(T_0(\infty, 0) \in [t, t+\epsilon) \mid T(\infty, 0) \geq t, V = \infty)} = \lim_{\epsilon \downarrow 0} \frac{P(T_1(s, 0) \in [t, t+\epsilon) \mid T(s, 0) \geq t, V = s)}{P(T_0(s, 0) \in [t, t+\epsilon) \mid T(s, 0) \geq t, V = s)}$$

Crucially, this assumption can allow for anticipation effects. This assumption presumes that the relative rates of vaccine-preventable to vaccine-irrelevant infection are stable for all strata where

participants are unvaccinated and uninfected at time t. Under this assumption, we can identify $\alpha_0(t)$ as follows.

$$\begin{aligned}\alpha_0(t) &= \lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty,0) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)}{P(T_0(\infty,0) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)} \\ &= \lim_{\epsilon \downarrow 0} \frac{P(T_1(s,0) \in [t, t+\epsilon) \mid T(s,0) \geq t, V = s)}{P(T_0(s,0) \in [t, t+\epsilon) \mid T(s,0) \geq t, V = s)} \qquad \text{s.t. } s > t \\ &= \lim_{\epsilon \downarrow 0} \frac{P(J = 1, T \in [t, t+\epsilon) \mid T \geq t, V > t)}{P(J = 0, T \in [t, t+\epsilon) \mid T \geq t, V > t)}\end{aligned}$$

Then using identical arguments as above, we obtain

$$\theta_C(t-v) = \frac{P(J = 1 \mid T = t, V = v)}{P(J = 0 \mid T = t, V = v)} \times \left(\frac{P(J = 1 \mid T = t, V > t)}{P(J = 0 \mid T = t, V > t)} \right)^{-1}$$

B.3.2 Sensitivity Analyses

We can parametrize violations of bias equivalence (Assumption 5) using a proportionality constant, $\beta(t)$.

$$\lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)}{P(T_1(\infty,t) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)} = \beta(t) \times \lim_{\epsilon \downarrow 0} \frac{P(T_0(\infty,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)}{P(T_0(\infty,t) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)}.$$

or equivalently,

$$\lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)}{P(T_0(\infty,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)} = \beta(t) \times \lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty,t) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)}{P(T_0(\infty,t) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)}.$$

Defining $\alpha_0(t)$ as the left hand side of the equation above, under Assumptions 6-7, we can show that

$$\begin{aligned}\alpha_0(t) &= \beta(t) \times \lim_{\epsilon \downarrow 0} \frac{P(T_1(\infty,t) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)}{P(T_0(\infty,t) \in [t, t+\epsilon) \mid T(\infty,0) \geq t, V = \infty)} \\ &= \beta(t) \times \lim_{\epsilon \downarrow 0} \frac{P(J = 1, T \in [t, t+\epsilon) \mid T \geq t, V = \infty)}{P(J = 0, T \in [t, t+\epsilon) \mid T \geq t, V = \infty)}\end{aligned}$$

Hence,

$$\begin{aligned}\theta_C(t-v) = \lim_{\epsilon \downarrow 0} & \frac{P(J = 1, T \in [t, t+\epsilon) | T \geq t, V = v)}{P(J = 0, T \in [t, t+\epsilon) | T \geq t, V = v)} \\ & \times \left(\beta(t) \times \lim_{\epsilon \downarrow 0} \frac{P(J = 1, T \in [t, t+\epsilon) \mid T \geq t, V = \infty)}{P(J = 0, T \in [t, t+\epsilon) \mid T \geq t, V = \infty)} \right)^{-1}\end{aligned}$$

$$= \frac{P(J=1 \mid T=t, V=v)/P(J=1 \mid T=t, V=\infty)}{P(J=0 \mid T=t, V=v)/P(J=0 \mid T=t, V=\infty)} \times (\beta(t))^{-1}$$

Hence, the causal VE parameter is identified by the conditional odds ratio of infection causes up to a sensitivity parameter.

Similarly, we can parametrize violations of exposure deferral irrelevance

$$\lim_{\epsilon \downarrow 0} \frac{P(T_j(v,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)}{P(T_j(v,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)} = \gamma(t)$$

Under this model, following the same arguments from our identification result in Section 1.3.1, we obtain

$$\begin{aligned}
\theta_C(t-v) = &= \lim_{\epsilon \downarrow 0} \frac{P(T_1(v,0) \in [t, t+\epsilon) | T(v,0) \geq t, V=v)}{P(T_0(v,t) \in [t, t+\epsilon) | T(v,0) \geq t, V=v)} \times [\alpha_0(t)]^{-1} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(T_1(v,0) \in [t, t+\epsilon) | T(v,0) \geq t, V=v)}{\gamma(t) \times P(T_0(v,0) \in [t, t+\epsilon) | T(v,0) \geq t, V=v)} \times [\alpha_0(t)]^{-1} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(T_1 \in [t, t+\epsilon) | T \geq t, V=v)}{P(T_0 \in [t, t+\epsilon) | T \geq t, V=v)} \times [\alpha_0(t)\gamma(t)]^{-1} \\
&= \lim_{\epsilon \downarrow 0} \frac{P(J=1, T \in [t, t+\epsilon) | T \geq t, V=v)}{P(J=0, T \in [t, t+\epsilon) | T \geq t, V=v)} \times [\alpha_0(t)\gamma(t)]^{-1}
\end{aligned}$$

And following the same arguments to identify $\alpha_0(t)$, we identify $\theta_C(t-v)$ up the multiplicative constant $\gamma(t)$.

$$\theta_C(t-v) = \frac{P(J=1 \mid T=t, V=v)/P(J=0 \mid T=t, V=v)}{P(J=1 \mid T=t, V=\infty)/P(J=0 \mid T=t, V=\infty)} \times [\gamma(t)]^{-1}$$

Finally, we can parametrize violations of the NCO assumption

$$P(T_0(v,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v) = \xi(t-v) \cdot P(T_0(\infty,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)$$

Following the same steps as the main identification result, we obtain

$$\theta_C(t-v) = \lim_{\epsilon \downarrow 0} \frac{P(J=1, T \in [t, t+\epsilon) \mid T \geq t, V=v)}{P(J=1, T \in [t, t+\epsilon) \mid T \geq t, V=v)} \times [\alpha_0(t)\xi(t-v)]^{-1}$$

And $\alpha_0(t)$ is identified as before, leading to modified version of the general identification result.

$$\theta_C(t-v) = \frac{P(J=1 \mid T=t, V=v)/P(J=0 \mid T=t, V=v)}{P(J=1 \mid T=t, V=\infty)/P(J=0 \mid T=t, V=\infty)} \times [\xi(t-v)]^{-1}$$

Taken together, a sensitivity analysis could consider

$$\theta_C(t-v) = \frac{P(J=1 \mid T=t, V=v)/P(J=0 \mid T=t, V=v)}{P(J=1 \mid T=t, V=\infty)/P(J=0 \mid T=t, V=\infty)} \times [\beta(t) \cdot \gamma(t) \cdot \xi(t-v)]^{-1}$$

Where $\beta(t)$ parametrizes the departure from bias equivalence, $\gamma(t)$ parametrizes the departure from exposure deferral irrelevance, and $\xi(t-v)$ parametrizes the departure from the NCO assumption on the hazard ratio scales.

B.3.3 Decompositions of the causal hazard ratio

Below, we introduce some decompositions of the causal hazard ratio. The first decomposition illustrates a potential cross-world depletion of susceptibles bias which complicates interpretation. The latter decomposition retains interpretability, and shows that the target parameter equals the product of a "pure" vaccine effect in a hypothetical world where exposure is deferred and a term that captures the impact of subinfectious exposures.

The first decomposition reduces the target parameter to the product of (i) the vaccine effect in the immediate exposure world and (ii) the effect of exposure deferral.

$$\begin{aligned} VE_C(t-v) &= 1 - \lim_{\epsilon \downarrow 0} \frac{P(T_1(v,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)}{P(T_1(\infty,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)} \\ &= 1 - \lim_{\epsilon \downarrow 0} \frac{P(T_1(v,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)}{P(T_1(\infty,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)} \times \frac{P(T_1(\infty,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)}{P(T_1(\infty,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)} \end{aligned}$$

While this decomposition is algebraically correct and conceptually appealing, we find it difficult to interpret due to the potential for "cross-world depletion of susceptibles". To illustrate, we focus on the first term of the decomposition.

$$\frac{P(T_1(v,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)}{P(T_1(\infty,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V=v)}$$

Focusing first on the numerator, all individuals satisfying $\{T(v,0) \geq t, V=v\}$ have a positive probability of experiencing the event $\{T_1(v,0) \in [t, t+\epsilon)\}$. Now focusing on the denominator, some individuals belonging to $\{T(v,0) \geq t, V=v\}$ may have experienced the event $\{T_1(\infty,0) < t\}$, and therefore would have zero probability of experiencing $\{T_1(\infty,0) \in [t, t+\epsilon)\}$. By applying law of

total probability, we can write the denominator as follows.

$$
\begin{aligned}
&P(T_1(\infty,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v) \\
&= P(T_1(\infty,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, T_1(\infty,0) \geq t, V = v) \times P(T_1(\infty,0) \geq t \mid T(v,0) \geq t, V = v) \\
&+ \underbrace{P(T_1(\infty,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, T_1(\infty,0) < t, V = v)}_{=0} \times P(T_1(\infty,0) < t \mid T(v,0) \geq t, V = v)
\end{aligned}
$$

Hence, the numerator describes participants in the principal stratum $S := \{T(v,0) \geq t, V = v\}$ while the denominator only describes a subset of participants in S who satisfy $S^* := S \cap \{T_1(\infty,0) \geq t\}$. If there is a high burden of infections and unvaccinated persons are particularly at risk of succumbing to infection before t, it is plausible that the distribution of risk factors among participants in S^* will be very different from those in S. This illustrates how cross-world comparisons can run into their own depletion of susceptibles issues, which complicates the interpretation of the first term as a meaningful notion of a "vaccine effect".

However, we believe the following decomposition of the causal estimand target has a clearer interpretation.

$$
\begin{aligned}
VE_C(t-v) &= 1 - \lim_{\epsilon \downarrow 0} \frac{P(T_1(v,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)}{P(T_1(\infty,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)} \\
&= 1 - \lim_{\epsilon \downarrow 0} \frac{P(T_1(v,0) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)}{P(T_1(v,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)} \times \frac{P(T_1(v,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)}{P(T_1(\infty,t) \in [t, t+\epsilon) \mid T(v,0) \geq t, V = v)}
\end{aligned}
$$

The first term on the right captures the subclinical infections effect – or the effect of delaying community exposure among vaccinated survivors – and the latter captures the "pure vaccination" effect among under hypothetical exposure deferral. When there is low probability of subclinical infections, the first term will be approximately 1 and the target estimand will approximately equal the pure vaccination effect in vaccinated survivors in a world where exposure was hypothetically deferred.

B.4 Proof of Proposition 1

A complete proof of the consistency and asymptotic normality of the semiparametric logistic regression model can be found in Chen [6]. We review the key elements of their proof.

B.4.1 Notation and conditions

Note that the semiparametric logistic regression model can be written as follows.

$$E[J|X] = \text{expit}\{\beta_0^T I(V \leq T)\underline{\psi}(T-V) - \alpha_0(T)\}$$

Where $\alpha_0(t)$ is an unspecified function defined on the positive real numbers and $\beta_0 \in \mathbb{R}^d$ is the Euclidean parameter of interest describing the time-varying VE. We use a finite-dimensional basis expansion of time-since-vaccination, $\{\psi_j(\cdot)\}_{j=1}^d$. The simplest possible specification of the basis expansion is a polynomial expansion of degree 1 ($d = 2$), meaning that $\psi_1(T-V) = 1, \psi_2(T-V) = T - V$, which assumes that VE changes linearly in time-since-vaccination on the log hazard/odds scale.

Suppose without loss of generality that $T \in [0, 1]$. The sieve estimator will approximate α_0 with a basis expansion that increases with the sample size. The approach will be to divide $[0, 1]$ into K equally spaced knots. Let $\underline{\phi}_K \equiv \{\phi_j(v, K)(\cdot)\}_{j=1}^M$ denote the collection of $M = K + v$ B-splines of degree v with K knots.

Suppose we approximate the true unknown function $\alpha_0(t)$ with the basis expansion associated with $\underline{\phi}_K$.

$$\alpha_0(t) \approx \sum_{j=1}^{M} \alpha_j \phi_j(t)$$

With $\underline{\alpha} \in \mathbb{R}^M$. Let $\theta_K(x)$ be the linear predictor associated with the logistic regression model with basis expansion for $\alpha_0(t)$ specified above.

$$\theta_K(t, v) = \beta^T I(v \leq t)\underline{\psi}(t - v) + \underline{\alpha}^T \underline{\phi}_K(t)$$

Moreover, define the following matrices

$$\mathcal{W} \in \mathbb{R}^n \times \mathbb{R}^d \text{ s.t. } \mathcal{W}_{ij} := I(V_i \leq T_i)\psi_j(T_i - V_i)$$
$$\mathcal{I} \in \mathbb{R}^n \times \mathbb{R}^M \text{ s.t. } \mathcal{I}_{ij} := \phi_{Kj}(T_i - V_i)$$

We define a Hölder class $\mathcal{G}(v_0, \gamma, c_0)$ as the collection of functions g such that

$$|g^{v_0}(x_1) - g^{(v_0)}(x_0)| \leq c_0|x_1 - x_0|^\gamma \quad \forall x_1, x_0 \in [0, 1]$$

We introduce the following conditions.

Condition S1. *Time-varying VE function lies in parametric class* Assume that the time-varying VE function $f(\cdot)$ lies in a function class spanned by the d-dimensional basis $\underline{\psi} := \{\psi_j\}_{j=1}^d$ with support on $\mathbb{R}^+$. In other words, for some $\beta_0 \in R^0 \subset \mathbb{R}^d$, where R is a bounded and convex set and R^0 is the interior of R.

$$f(t-v) = \beta_0^T \underline{\psi}(t-v)$$

Condition S2. *Nuisance parameter lies in Hölder Class* Suppose that $\alpha_0(t)$ lies in the class of functions $\mathcal{G}(v_0, \gamma, c_0)$ for $\gamma \in [0,1]$ consisting of functions satisfying the Hölder condition.

$$|g^{(v_0)}(x_1) - g^{(v_0)}(x_0)| \leq c_0 |x_1 - x_0|^\gamma \quad \forall\, x_1, x_0 \in [0,1]$$

Where c_0 is a fixed positive constant. Some common choices are Lipschitz functions $(v_0, \gamma) = (0,1)$, functions with bounded derivatives $(v_0, \gamma) = (1,0)$, and functions with bounded higher-order derivatives $(v_0, \gamma) = (r, 0)$. Define the smoothness index $p = v_0 + \gamma$. Assume the $\alpha_0(t)$ is also uniformly bounded.

$$\sup_{t \in [0,1]} |\alpha_0(t)| < \infty$$

Condition S3. *Nuisance basis expansion* Let $\underline{\phi}_K := \{\phi_j(\cdot; K, v)\}_{j=1}^M$ refer to collection of $M = K + v$ B-splines of degree v on $[0,1]$ with the following knot placements at j/K for $1 \leq j < K$. Let $S(v, K)$ denote the vector space of dimension M spanned by the basis.

Condition S4. *Positive support* For any $a \in \mathbb{R}^d$, $a \neq 0$, and $c \in \mathbb{R}$,

$$P(a^T I(V \leq T) \underline{\psi}(T - V) \neq c) > 0$$

This condition is required to ensure that the distribution of the linear predictor does not concentrate in a lower-dimensional hyperplane, ensuring β_0 is identifiable.

Condition S5. *Size of sieve space* Suppose $K_n = O_p(n^\lambda)$ for some $0 < \lambda$ and $K_n/n = o_p(1)$.

B.4.2 Proof of Proposition 1

Recall that $(\hat{\beta}_K, \hat{\alpha}_K)$ is a solution to a series of likelihood equations associated with the semiparametric logistic regression model with linear predictor θ_K. Using the Taylor series expansion of the log likelihood, we can write the score with respect to each parameter as follows

$$\left.\frac{\partial \ell_n(\beta,\alpha)}{\partial \beta_j}\right|_{\substack{\beta=\beta_{K0}\\ \alpha=\alpha_{K0}}} = \sum_{i=1}^{n} \text{expit}(\theta_K^*)[\mathcal{W}_i^T(\hat{\beta}_K - \hat{\beta}_{K0}) + \mathcal{I}_i^T(\hat{\alpha}_K - \hat{\alpha}_{K0})]I(v_i \le t_i)\psi_j(t_i - v_i)$$

$$\left.\frac{\partial \ell_n(\beta,\alpha)}{\partial \alpha_k}\right|_{\substack{\beta=\beta_{K0}\\ \alpha=\alpha_{K0}}} = \sum_{i=1}^{n} \text{expit}(\theta_K^*)[\mathcal{W}_i^T(\hat{\beta}_K - \hat{\beta}_{K0}) + \mathcal{I}_i^T(\hat{\alpha}_K - \hat{\alpha}_{K0})]\phi_k(t_i)$$

Where $\theta_K^* = (1-\lambda_n)\hat{\theta}_K + \lambda_n\theta_{K0}$ for some $\lambda_n \in [0,1]$. Set $\underline{Y} = (J_1, \ldots, J_n)^T$ and $\underline{D}_n = (-\text{expit}(\theta_{K0}(\underline{x}_1)), \ldots, -\text{expit}(\theta_{K0}(\underline{x}_n)))^T$, and let $\mathcal{A}_n$ be a diagonal matrix with iith entry $\exp\{\theta_k^*(x_i)\}/(1+\exp\{\theta_k^*(x_i)\})^2$. Define $\bar{\mathcal{W}}^T = \mathcal{W}^T\mathcal{A}_n^{1/2}$ and $\bar{\mathcal{I}} = \mathcal{I}^T\mathcal{A}_n^{1/2}$. We can rewrite the score equations as follows.

$$\begin{pmatrix} \bar{\mathcal{W}}^T\bar{\mathcal{W}} & \bar{\mathcal{W}}^T\bar{\mathcal{I}} \\ \bar{\mathcal{I}}^T\bar{\mathcal{W}} & \bar{\mathcal{I}}^T\bar{\mathcal{I}} \end{pmatrix} \begin{pmatrix} \hat{\beta}_K - \beta_{K0} \\ \hat{\alpha}_K - \alpha_{K0} \end{pmatrix} = \begin{pmatrix} \bar{\mathcal{W}}^T \\ \bar{\mathcal{I}}^T \end{pmatrix} \mathcal{A}_n^{-1/2}[\underline{Y} + \underline{D}_n]$$

Applying the Frisch–Waugh–Lovell theorem, we can obtain an expression in terms of the parameter of interest, $(\hat{\beta}_K - \beta_{K0})$, without dependence on α.

$$(\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{1/2}\mathcal{W})(\hat{\beta}_K - \beta_{K0}) = \mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{-1/2}(\underline{Y} - \underline{D}_n)$$

Where $\mathbf{P}_n = \mathcal{A}_n^{1/2}\mathcal{I}(\mathcal{I}^T\mathcal{A}_n\mathcal{I})^{-1}\mathcal{I}^T\mathcal{A}_n^{1/2}$. We recognize that each element of $\underline{Y} - \underline{D}_n$ is a residual of the form $Y_i - \text{expit}(\theta_{K0}(X_i)) = Y_i - \text{expit}(\theta_0(X_i)) + \text{expit}(\theta_0(X_i)) - \text{expit}(\theta_{K0}(X_i))$. Let the vector $\underline{\epsilon}_{Kn}$ capture the random error (with elements $Y_i - \text{expit}(\theta_0(X_i))$) and let the vector $\underline{b}_n$ reflect the approximation error (with elements $\text{expit}(\theta_0(X_i)) - \text{expit}(\theta_{K0}(X_i))$).

$$(\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{1/2}\mathcal{W})(\hat{\beta}_K - \beta_{K0}) = \mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{-1/2}(\underline{\epsilon}_{Kn} + \underline{b}_n)$$

We now note that the asymptotic behavior of $\sqrt{n}(\hat{\beta}_K - \beta_{K0})$ depends on three terms.

(C1) $n^{-1}\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{1/2}\mathcal{W}$

(C2) $n^{-1/2}\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{-1/2}\underline{b}_n$

(C3) $n^{-1/2}\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{-1/2}\underline{\epsilon}_{Kn}$

By Lemma 9 in Chen[6], the first term converges in probability to a matrix.

$$n^{-1}\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{1/2}\mathcal{W} \xrightarrow{p} \mathbf{\Sigma}$$

In Section 6.3.1 of Chen[6], we can establish a stochastic upper bound on the second term.

$$|n^{-1/2}\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{-1/2}\underline{b}_n| = O_P(n^{-1/2}MK^{-p}) + O_p(K^{-p})$$

The third term is of particular interest. We note that for any $\underline{a}^T \in \mathbb{R}^d$, we write

$$\underline{a}^T\mathcal{W}^T\mathcal{A}_0^{1/2}(\mathbf{I}-\mathbf{P}_0)\mathcal{A}_0^{-1/2}\underline{\epsilon}_{Kn} = \underline{a}^T\mathcal{W}^T\mathcal{A}_0^{1/2}(\mathbf{I}-\mathbf{P}_0)\underline{e}_n$$

Where e_n has elements $e_i = \epsilon_i/\sigma_i$ where $\epsilon_i = J_i - \text{expit}(\theta_0(X_i))$ and $\sigma_i = [\mathcal{A}_0^{1/2}]_{ii}$. We can write $n^{-1/2}\underline{a}^T\mathcal{W}^T\mathcal{A}_0^{1/2}(\mathbf{I}-\mathbf{P}_0)\underline{e}_n$ as $\sum_{j=1}^n c_j e_j$ where $E[e_j] = 0$, $\text{Var}(e_j) = 1$, and $E[e_j^4] < \infty$. By Lemma 9(a) in Chen[6], $\sum_{j=1}^n c_j^2 = O_p(1)$.

Defining $\underline{a}^T\mathcal{W}^T\mathcal{A}_0^{1/2} = (g_1, \ldots, g_n)$, it follows from Conditions 2 and 4 that $\max_{1\le i\le n}|g_i| = O_p(1)$ and $\sum_{i=1}^n g_i^2 = O_p(n)$. Let p_{ij} denote the (ij)-th element of $\mathbf{P}_0$. It follows from Lemma 8 of Chen[6] that $\sum_{i=1}^n p_{ij}^2 = p_{ii} = O_p(n^{-1}M)$. Hence,

$$\begin{aligned}
\sum_{j=1}^n |c_j|^3 &\le \max_{1\le j\le n}|c_j| \sum_{j=1}^n c_j^2 \\
&\le n^{-1/2} \max_{1\le j\le n}\left[|g_j| + \left(\sum_{i=1}^n g_i^2 \sum_{i=1}^n p_{ij}^2\right)^{1/2}\right] \sum_{j=1}^n c_j^2 \\
&= n^{-1/2}\left[O(1) + (O(n)O_p(n^{-1}M))^{1/2}\right] O_p(1) \\
&= o_p(1)
\end{aligned}$$

The first inequality is implied by the maximum being an upper bound. The second inequality comes from finding an upper bound on c_j using the triangle inequality. The penultimate equality comes from substituting in the rates for each step, and the final equality comes from manipulating the stochastic order notation.

Hence, by the Cramér-Wald device and the Liapounov central limit theorem, we obtain

$$n^{-1/2}\mathcal{W}^T\mathcal{A}_0^{1/2}(\mathbf{I}-\mathbf{P}_0)\underline{e}_n \rightsquigarrow N(0, \Sigma)$$

And equivalently

$$n^{-1/2}\mathcal{W}^T\mathcal{A}_0^{1/2}(\mathbf{I}-\mathbf{P}_0)\mathcal{A}_0^{-1/2}\underline{\epsilon}_{Kn} \rightsquigarrow N(0,\Sigma)$$

Rewinding to the start of the proof, recall

$$\underbrace{n^{-1}(\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{1/2}\mathcal{W})}_{\xrightarrow{p}\mathbf{\Sigma}}\sqrt{n}(\hat{\beta}_K-\beta_{K0}) = \underbrace{n^{-1/2}\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{-1/2}\underline{\epsilon}_{Kn}}_{\star} + \underbrace{n^{-1/2}\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{-1/2}\underline{b}_n}_{O_p(n^{-1/2}MK^{-p})+O_p(K^{-p})}$$

By Equation (21) in Chen[6],

$$\begin{aligned}
&\mathcal{W}^T\mathcal{A}_n^{1/2}(\mathbf{I}-\mathbf{P}_n)\mathcal{A}_n^{1/2}\underline{\epsilon}_K \\
&= \mathcal{W}^T\mathcal{A}_0^{1/2}(\mathbf{I}-\mathbf{P}_0)\mathcal{A}_0^{-1/2}\underline{\epsilon}_K + \mathcal{W}^T(\mathcal{A}_0-\mathcal{A}_n)\mathcal{A}_0^{-1/2}\mathbf{P}_0\mathcal{A}_0^{-1/2}\underline{\epsilon}_K \\
&+ \mathcal{W}^T\mathcal{A}_n\mathcal{I}\left[(\mathcal{I}^T\mathcal{A}_0\mathcal{I})^{-1}-(\mathcal{I}^T\mathcal{A}_n\mathcal{I})^{-1}\right]\mathcal{I}^T\underline{\epsilon}_K
\end{aligned}$$

The first term on the right hand side looks familiar from the previous CLT result. The second term can be stochastically bounded using Proposition 1 from Chen[6].

$$\mathcal{W}^T(\mathcal{A}_0-\mathcal{A}_n)\mathcal{A}_0^{-1/2}\mathbf{P}_0\mathcal{A}_0^{-1/2}\underline{\epsilon}_K = O_P((nK^{-2p}M)^{1/2}) + O_P(M(\log(n)^{3/2})$$

The third term can be stochastically bounded using Proposition 2 from Chen[6].

$$\mathcal{W}^T\mathcal{A}_n\mathcal{I}\left[(\mathcal{I}^T\mathcal{A}_0\mathcal{I})^{-1}-(\mathcal{I}^T\mathcal{A}_n\mathcal{I})^{-1}\right]\mathcal{I}^T\underline{\epsilon}_K = O_P(n^{-1/2}M^{5/2}\log(n)) + O_P(M^{3/2})$$

Hence, in order for $\star \rightsquigarrow N(0,\mathbf{\Sigma})$, we require that each of the stochastically bounded terms go to zero. Recall that $M \approx K \approx n^{\lambda}$.

(C1) $\lim_{n\to\infty} n^{-1/2}MK^{-p} = \lim_{n\to\infty} n^{\lambda(1-p)-1/2} \implies \lambda(1-p) < 1/2$

(C2) $\lim_{n\to\infty} n^{-1/2}(nK^{-2p}M)^{1/2} = \lim_{n\to\infty} n^{\frac{\lambda(1-2p)}{2}} \implies \lambda(1-2p) < 0 \implies p > 1/2$

(C3) $\lim_{n\to\infty} n^{-1/2}M^{3/2} = \lim_{n\to\infty} n^{3\lambda/2-1/2} \implies \lambda < 1/3$

(C4) $\lim_{n\to\infty} n^{-1/2}M\log\{n\}^{3/2} = \lim_{n\to\infty} n^{\lambda-1/2} \implies \lambda < 1/2$

(C5) $\lim_{n\to\infty} n^{-1/2}(n^{-1/2}M^{5/2}\log\{n\}) = \lim_{n\to\infty} n^{-1+5\lambda/2} \implies \lambda < 2/5$

(C6) $\lim_{n\to\infty} K^{-p} = \lim_{n\to\infty} n^{-\lambda p} \implies \lambda, p > 0$

Taken together, we conclude that $0 < \lambda < 1/3$ and $p > 1/2$. When this occurs, the second and third terms on the right hand side of (21) by Chen[6] go to zero. Hence,

$$\sqrt{n}(\hat{\beta}_K - \beta_{K0}) \rightsquigarrow N(0, \mathbf{\Sigma}^{-1})$$

Moreover, from Theorem 1(b) in Chen[6], $|\beta_{K0} - \beta_0| = O(K^{-p})$. Hence,

$$\sqrt{n}(\hat{\beta}_K - \beta_0) = \underbrace{\sqrt{n}(\hat{\beta}_K - \beta_{K0})}_{\rightsquigarrow N(0,\mathbf{\Sigma}^{-1})} + \underbrace{\sqrt{n}(\beta_{K0} - \beta_0)}_{O(n^{-\lambda p+1/2})}$$

Hence, under the additional undersmoothing condition $\lambda p > 1/2 \equiv p > 1/(2\lambda)$, the latter term is $o_P(1)$, which resolves the asymptotic distribution of $\sqrt{n}(\hat{\beta}_K - \beta_0)$.

B.5 Proof of Proposition 2

Suppose our data of size n are sampled independently and identically from $P_0 \sim M$, where M is the model containing distributions satisfying the following conditional mean equality.

$$Q_0(V,T) \equiv P(J = 1|V,T) = \text{expit}(\beta_0^T I(V \leq T) \cdot \underline{\psi}(T - V) + \alpha_0(T)) \equiv \text{expit}(\beta_0^T \underline{Z} + \alpha_0(T))$$

For some $\beta_0 \in \mathbb{R}^d$ and some function α_0. In some cases, we may use $\underline{Z} := I(V \leq T) \cdot \underline{\psi}(T - V)$ as shorthand to denote the derived regressors. We will pursue estimation and inference on β_0 using a *targeted maximum likelihood estimation (TMLE) procedure*.

The TMLE algorithm first obtains initial estimates of the relevant portions of the distribution of the observed data. Then the initial fits are updated to optimize the bias-variance tradeoff. This is operationalized by updating the initial fits using a logistic working model that solves the efficient score equation – defined as the efficient influence function – in the semiparametric model defined above through the use of a "clever covariate".

B.5.1 Derivation of EIC

A key component of our TMLE algorithm is defining the efficient influence function/efficient influence curve of β_0 in the semiparametric partially-linear logistic regression model. To identify the efficient influence curve, we note that an influence function that is orthogonal to the nuisance tangent space (i.e., scores obtained by fluctuation α_0) is the unique EIF. Hence, a good place to

start is identifying the form of the nuisance tangent space. The nuisance tangent space will take the following form.

$$T_{nuis,\alpha_0}(P_0) = \{h(T)(J - Q_0(V,T)) : h \in L_2\}$$

For arbitrary squared integrable function h. We wish to construct a path $Q(\epsilon)$ through Q such that the score at $\epsilon = 0$ is orthogonal to the nuisance tangent space and is hence the EIC. Consider the candidate paths that fluctuate $\alpha_0(t)$ in the direction indexed by h_1 and fluctuate β_0 in the direction b.

$$Q_{0,b,h_1}(\epsilon) = \text{expit}((\beta_0^T + \epsilon \cdot b^T)\underline{Z} + \alpha_0(T) + \epsilon h_1(T))$$

The score of this path at $\epsilon = 0$ is given by

$$\underline{S}(h_1, \underline{b}) \equiv \left(\underline{b}^T \underline{Z} + h_1(T)\right)(J - Q_0(A,V,T))$$

Now we need to identify the forms of h_1 and b such that orthogonality holds, meaning for any $h(T)$, we have

$$\begin{aligned}
0 &= \mathbb{E}[S(h_1, b) \times h(T)(J - Q_0(A,V,T))] \\
&= \mathbb{E}[\left(\underline{b}^T \underline{Z} + h_1(T)\right)(J - Q_0(A,V,T)) \times h(T)(J - Q_0(A,V,T))] \\
&= \mathbb{E}[\left(\underline{b}^T \underline{Z} + h_1(T)\right) \times h(T) \times (J - Q_0(A,V,T))^2] \\
&= \mathbb{E}[\left(\underline{b}^T \underline{Z} + h_1(T)\right) \times h(T) \times Q_0(A,V,T)(1 - Q_0(A,V,T))]]] \\
&= \mathbb{E}[\left(\underline{b}^T I(V \leq T)\underline{\psi}(T - V) + h_1(T)\right) \times h(T) \times Q_0(A,V,T)(1 - Q_0(A,V,T))]
\end{aligned}$$

Where the penultimate equality follows from Tower law with respect to J. The unique solution in h_1 is given by the following.

$$h_1^*(Q_0)(t) = -\underline{b}^T \frac{\mathbb{E}_V[I(V \leq T)\underline{\psi}(T - V)Q_0(1 - Q_0)(V,T)|T = t]}{\mathbb{E}_V[Q_0(1 - Q_0)(V,T)|T = t]}$$

Plugging back into the equation for the score, we obtain

$$\underline{S}(h_1^*, \underline{b}) = \underline{b}^T \left(I(V \leq T)\underline{\psi}(T - V) - \frac{\mathbb{E}_V[I(V \leq T)\underline{\psi}(T - V)Q_0(1 - Q_0)(V,T)|T]}{\mathbb{E}_V[Q_0(1 - Q_0)(V,T)|T]}\right)(J - Q_0(V,T))$$

So for any arbitrary direction $\underline{b}$, the efficient directional score is given above. This implies that the EIF (or efficient score vector) is proportional to the following.

$$D(Q_0)(J,V,T) \propto (J - Q_0(V,T)) \left(I(V \leq T)\underline{\psi}(T-V) - r_0(T) \right)$$
$$r_0(T) := \frac{\mathbb{E}_V[I(V \leq T)\underline{\psi}(T-V)Q_0(1-Q_0)(V,T)|T]}{\mathbb{E}_V[Q_0(1-Q_0)(V,T)|T]}$$

To correctly identify the EIF, we must scale by the information matrix, defined below as follows.

$$\begin{aligned}
I_{eff} &:= \mathbb{E}[D(Q_0)D(Q_0)^T] \\
&= \mathbb{E}\left[(\underline{Z} - r_0(T)) (\underline{Z} - r_0(T))^T (J - Q_0(V,T))^2 \right] \\
&= \mathbb{E}_T \left[\mathbb{E}_V \left[(\underline{Z} - r_0(T)) (\underline{Z} - r_0(T))^T Q_0(V,T) (1 - Q_0(V,T)) \Big| T \right] \right] \\
&= \mathbb{E}_T \left[\mathbb{E}_V \left[\underline{Z}\underline{Z}^T Q_0(V,T) (1 - Q_0(V,T)) \Big| T \right] - r_0(T) r_0(T)^T \mathbb{E}_V (Q_0(V,T) (1 - Q_0(V,T)) |T) \right] \\
&= \mathbb{E}_T \left[\mathbb{E}_V (Q_0(V,T) (1 - Q_0(V,T)) |T) \left(\frac{\mathbb{E}_V \left[\underline{Z}\underline{Z}^T Q_0(V,T) (1 - Q_0(V,T)) \Big| T \right]}{\mathbb{E}_V (Q_0(V,T) (1 - Q_0(V,T)) |T)} - r_0(T) r_0(T)^T \right) \right]
\end{aligned}$$

Where the first equality is definitional, the second substitutes from above, the third follows by iterated expectation, the fourth from reducing the cross terms, and the fifth from using a multiply divide trick.

Hence, the EIF is as follows.

$$D(Q_0) = [I_{eff}]^{-1} \left(I(V \leq T)\underline{\psi}(T-V) - r_0(T) \right) (J - Q_0(V,T))$$

B.5.2 Description of TMLE estimator

Now that the EIC is derived, we can compute the TMLE estimator as follows. Compute initial (i.e. flexible) estimators of Q_0 (or equivalently, β_0 and $\alpha(t)$) using flexible models such as generalized additive model or random forest. Let Q^0 be defined as your initial estimate and let $p_i^0 = \text{expit}(I(V_i \leq T_i)\underline{\psi}(T_i - V_i)\beta^{(0)} + \alpha^{(0)}(T_i))$. Then for each iteration $k = 0, 1, 2, \ldots$, estimate the following quantities.

$$w_i^{(k)} = p_i^{(k)}(1 - p_i^{(k)})$$
$$r^{(k)}(t) = \frac{\mathbb{E}_V[I(V_i \leq T_i)\underline{\psi}(T_i - V_i)w_i^{(k)}|T = t]}{\mathbb{E}_V[w_i^{(k)}|T = t]}$$

The latter quantity can be estimated by kernel smoothing, local regression, or fitting a GAM for each column of $I(V_i \leq T_i)\underline{\psi}(T_i - V_i)$ with weight w. Let $r_i^{(k)} := r^{(k)}(T_i)$ denote the vector at each T_i.

Then, compute the clever covariate $H_i^{(k)}$ which is a $n \times d$ matrix.

$$H_i^{(k)} := I(V_i \leq T_i)\underline{\psi}(T_i - V_i) - r^{(k)}(T_i)$$

Next, we regress J on the clever covariate with the previous iteration's probabilities as an offset to solve for the update $\epsilon^{(k)}$.

$$\text{logit}(p_i^{(new)}) = \text{logit}(p_i^{(k)}) + \epsilon^T H_i^{(k)}$$

After updating the fit, we update β and $\alpha(t)$ as follows.

$$\beta^{(k+1)} \leftarrow \beta^{(k)} + \epsilon^{(k)}$$
$$\alpha^{(k+1)}(t) \leftarrow \text{smooth}_t\{\text{logit}(p_i^{(new)}) - I(V_i \leq T_i)\underline{\psi}(T_i - V_i)\beta^{(k+1)}\}$$

Where smooth_t refers to fitting a smoother over t using smoothing spline or a GAM. This represents a profile smoothing update to $\alpha^{(k)}$, as $\beta^{(k+1)}$ is treated as known.

Finally, we recompute the event probabilities.

$$p_i^{(k+1)} = \text{expit}(I(V_i \leq T_i)\underline{\psi}(T_i - V_i)^T\beta^{(k+1)} + \alpha^{(k+1)}(T_i))$$

This iterative process is repeated until the magnitude of the nuisance parameter updates, ϵ, is small. Specifically, we repeat until $||\epsilon^{(k)}|| < \text{tol}$ where tol is a tolerance level (e.g., 10^{-6}). Our final estimate $\hat{\beta}_{TMLE}$ is obtained when the tolerance criterion is reached, and let p_i^{final} refer to the final targeted fitted probabilities.

The variance of the $\hat{\beta}_{TMLE}$ is estimated by the n^{-1}-scaled empirical variance fo the EIF,

$$\widehat{\text{Var}}(\hat{\beta}_{TMLE}) \approx n^{-1}\left(\frac{1}{n}\sum_{i=1}^{n}\hat{D}_{P_n^*}(O_i)\hat{D}_{P_n^*}(O_i)^T\right)$$

Where $\hat{D}_{P_n^*}(O_i) = \hat{I}_{eff}^{-1}(J_i - p_i^{final})H_i^{(final)}$ and $\hat{I}_{eff} = P_n[(J - p^{(final)})H^{(final)}H^{(final)T}]$.

B.5.3 Consistency and Asymptotic normality

Let P_n^* be the targeted distribution for $\beta(P)$ obtained after updating the initial estimator $P_{n,0} \in \mathcal{M}$. Denote $\tilde{r}_n = \tilde{r}_{P_{n,0}}$ and $Q_n^* := Q_{P_n^*}$. Following the results by van der Laan and Gilbert [7], we propose the following conditions which $\beta(P_n^*)$ is asymptotically linear and efficient for $\beta(P)$.

Condition S6. *Uniformly bounded and invertible* $\underline{\psi}(T-V)$ if P-uniformly bounded and $\mathbb{E}_P[\underline{\psi}(T-V)\underline{\psi}(T-V)^T]$ is an invertible matrix.

Condition S7. *Positivity* There exists some scalar $\delta > 0$ such that $\mathbb{P}(1-\delta > Q_n^*(V,T) > \delta) = P(1-\delta > Q_0(V,T) > \delta) = 1$.

Condition S8. *Donsker* The estimators $\tilde{r}_n$ and Q_n^* fall with probability one in a uniformly bounded P-Donsker function class.

Note that Condition S8 is not strictly necessary, and can be relaxed to include more flexible machine learning estimators that use sample-splitting and cross-fitting [8,9].

Condition S9. *Sufficient nuisance convergence rates* $||\tilde{r}_n - r_0|| = o_P(n^{-1/4})$ and $||Q_n^* - Q_0|| = o_P(n^{-1/4})$ where $||\cdot||$ refers to the L_P^2 norm.

The proof of the CAN of $\beta_{TMLE} := \beta(P_n^*)$ is standard and we refer to van der Laan and Rose [8] for details. It is based on writing the difference between the targeted estimator and targeted parameter in three terms.

$$\beta(P_n^*) - \beta(P) = \underbrace{-P_n D_{P_n^*}}_{1} + \underbrace{(P_n - P)D_{P_n^*}}_{2} + \underbrace{\{\beta(P_n^*) - \beta(P) + PD_{P_n^*}}_{3})$$

The first term is $o_P(n^{-1/2})$ by construction of the TMLE, since we choose our estimate of the relevant components of the data generating distribution to solve the efficient score equation (at least at $o_P(n^{-1/2})$ rate). The third term is $o_P(n^{-1/2})$ based on Lemma 1 of van der Laan and Gilbert [7]. Therefore, the expansion takes the form

$$\begin{aligned}\beta(P_n^*) - \beta(P) &= (P_n - P)D_{P_n^*} + o_P(n^{-1/2}) \\ &= (P_n - P)D_P + (P_n - P)\{D_{P_n^*} - D_P\} + o_P(n^{-1/2})\end{aligned}$$

Under Condition S7, D_P is uniformly bounded with finite variance. By the CLT, $\sqrt{n}(P_n - P)D_P$

goes to a normally distributed variable with the desired covariance structure. By Slutsky's theorem, what remains is to show

$$(P_n - P)\{D_{P_n^*} - D_P\} = o_P(n^{-1/2})$$

This empirical process term is negligible when $D_{P_n^*} - D_P$ falls in a Donsker class with probability 1. To establish this, if the nuisance estimators lie in P-Donsker classes (Condition S8) and $\tilde{r}_n \to r_0$ and $Q_n^* \to Q_0$ converge in $L^2(P)$ at appropriate rates (Condition S9), then the class $\{D(\cdot, Q, r)\}$ is P-Donsker. Because the Donsker class is invariant to Lipschitz transformations, $D_{P_n^*} - D_P$ lies in a Donsker class. Stochastic equicontinuity of empirical processes over Donsker classes gives $(P_n - P)\{D_{P_n^*} - D_P\} = o_P(n^{-1/2})$ as desired.

B.6 Example Data Application

Consider a hypothetical study that enrolls $n = 10,000$ individuals at a given calendar date and follows them for 1 year. Suppose that 85% of study participants are vaccinated at some point during follow-up and 15% are not.

Suppose that participant-level hazard of experiencing infection are determined by an unmeasured frailty variable, U_i. Suppose $U_i \sim \text{logNormal}(0, \sigma_U = 2.5)$ consistent with a very heterogeneous frailty distribution. Suppose that vaccination date is drawn independently from $V \sim \text{Unif}(0, 1.15)$ and any vaccination date greater than 1 is treated as never vaccinated. There was no association between vaccination date and no post-vaccination disinhibition. We assumed

We assume the true hazards functions for the vaccine-preventable and vaccine-irrelevant infections are as follows

$$h_1(t) = h_{01}(t) \exp(I(V_i \leq t)[\beta_0 + \beta_1(t - V_i) + U_{i1} - U_{i0}] + U_{i0})$$
$$h_0(t) = h_{00}(t) \exp(I(V_i \leq t)[U_{i1} - U_{i0}] + U_{i0})$$

Where $h_0^Y(t) = 0.16 \times (1 + 0.5 \sin(t * 2\pi))$ and $h_0^N(t) = 0.08$. Hence, the hazard of COVID-19 fluctuates sinusoidally in calendar time while the hazard of off-target infection is constant, implying a sinusoidal change-of-scale parameter $\alpha_0(t)$. The vaccine has no effect on the hazard of vaccine-irrelevant infection but does affect the hazard of vaccine-preventable infections. The time-varying VE curve $f(\tau) = \beta_0 + \beta_1 \tau$ is linear in time-since-vaccination on the log hazard scale. We consider

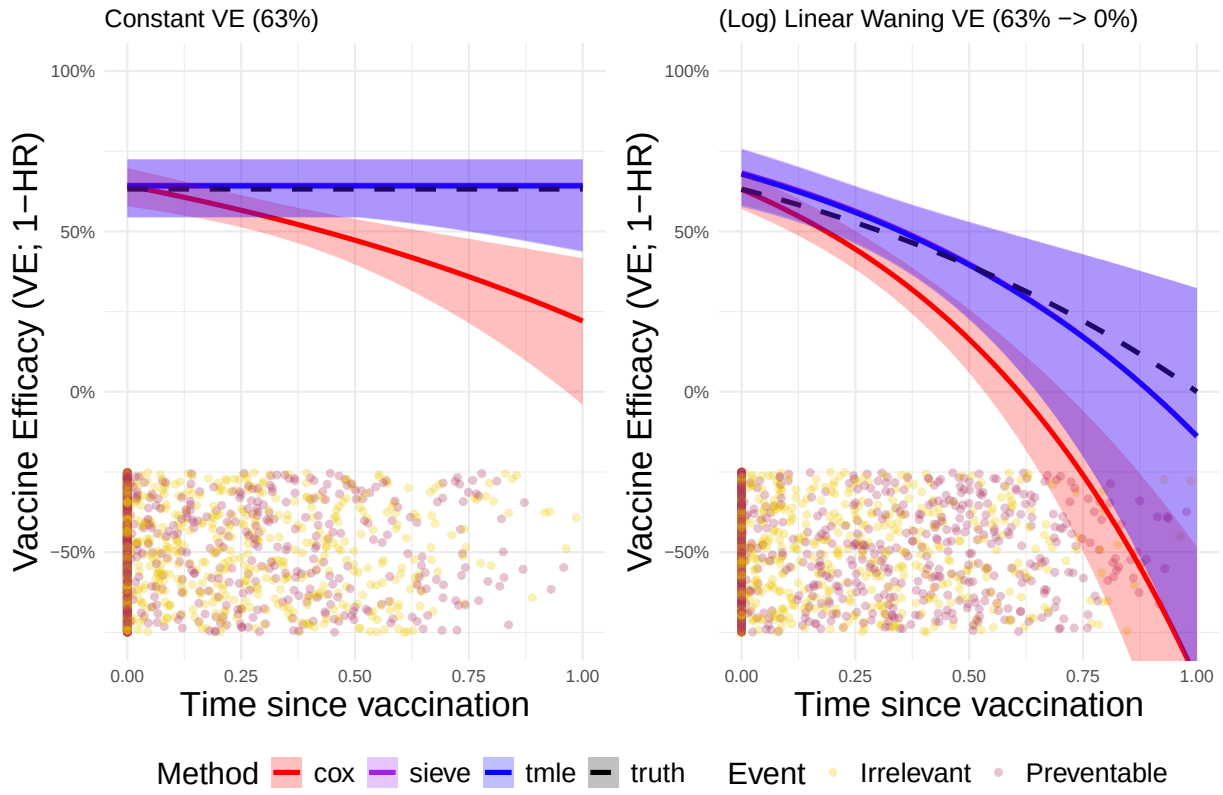

Figure B.1: Results of example data application comparing naive Cox regression to our proposed estimators. Sample size $n = 10,000$, emulating an observational study where there is hidden association between early vaccination and higher susceptibility to infection and participants exhibit behavioral disinhibition and increased pathogen exposure post-vaccination.

two scenarios. Scenario 1 assumes constant, moderate VE over time ($\beta_0 = -1, \beta_1 = 0$) while scenario 2 assumes a moderate initial VE that wanes linearly on the log-hazard scale to zero after one year ($\beta_0 = -1, \beta_1 = 1$).

In our example data application, we will compare (i) naive Cox regression, (ii) Sieve estimation approach, and (iii) TMLE approach. We assume correctly specified linear models for the VE waning model in each case. We applied a monotone non-decreasing constraint to the parameter estimates obtaind using the sieve and TMLE methods, and obtained associated 95% confidence intervals by projecting monte carlo samples from the unconstrained asymptotic distribution onto the constrained space.

We show our results in B.1. As anticipated, the standard Cox regression estimator of the time-varying VE curve exhibited exaggerated VE waning due to depletion of susceptibles bias. Qualitatively, the Cox regression model incorrectly inferred that VE was waning even in the setting where VE was constant. In contrast, both the TMLE and sieve estimation approaches which leveraged the vaccine-irrelevant infections both exhibited low bias and correctly identified the true time-varying VE shape under both waning and constant scenarios. Moreover, the confidence intervals for the TMLE and sieve estimators both covered the true VE function well. As anticipated, the confidence intervals for the TMLE and sieve estimators were wider than for cox regression, since they did not make use of censored participants in their estimator.

B.7 Numerical Experiments

The goals of our simulations were twofold (i) to confirm the asymptotic properties of our proposed estimators and (ii) to compare the performance of our proposed estimators to a standard Cox regression analysis which ignored vaccine irrelevant infections.

We evaluated performance across two settings: a constant VE setting where vaccine efficacy was constant over time at 63% ($f(\tau) = -1$) and a linear VE waning setting where efficacy waned linearly on the log hazard scale from 63% immediately after vaccination to 0% after 1 year ($f(\tau) = -1 + \tau$). In all simulations, we assumed that the baseline hazard of off-target events was constant over time $\lambda_{00}(t) = 0.03$ but the baseline hazard for vaccine-preventable infections followed a sinusoidal pattern, $\lambda_{01}(t) = \lambda_{01}(1 - 0.5\sin(2\pi t))$, meaning $\alpha_0(t)$ was sinusoidal. All our simulated studies were administrative censored at 1 year and contained 10,000 participants. We assumed that on average, approximately 85% of participants received the vaccine and 15% of participants remained unvaccinated by simulating vaccination date $V \sim \text{Unif}(0, 1.15)$. If the simulated vaccination date was greater than 1, they were considered unvaccinated.

We varied several different parameters that could challenge our ability to reliably estimate time-varying VE estimates.

(C1) Unmeasured heterogeneity in infection risk: we simulated frailties $U_i \sim \text{lognormal}(0, \sigma_U)$ and included them as unknown multiplicative factors in the hazard of both vaccine-preventable and vaccine-irrelevant infections.

(C2) Background vaccine-preventable infection intensity: prior work has suggested that depletion of susceptibles bias can be magnified during periods of high burden of infection. Hence, we varied the vaccine preventable infection burden by modulating $\lambda_{01} \in \{0.03, 0.06, 0.12\}$ reflecting low, moderate, and high infection burdens.

(C3) Behavioral disinhibition: to reflect changes in behavior influencing infection risk after vaccination, (U_0, U_1) from a standard bivariate log-normal distribution with marginal distributions lognormal$(0, \sigma_U = 1)$ and lognormal$(1, \sigma_U = 1)$ and weak correlation $\rho = 0.2$ The associated multiplicative contribution to the hazards is then $I(V \leq t)U_0 + I(V > t)U_1$, where the mean shift from U_0 to U_1 reflects the case where average participant risk increases after vaccination.

(C4) Non-random booster assignment: when booster assignment was random, the frailties U and vaccination date V were simulated separately from normal and uniform distributions. To accommodate a scenario where high-risk, susceptible individuals are prioritized for early vaccination, we simulated (U, V) from a bivariate normal copula model with marginal distributions lognormal$(0, 1)$ and Unif$(0, 1.15)$ with a negative correlation parameter $\rho = -0.7$. This copula model ensured that early vaccination date was associated with higher values of the frailty.

We considered three different estimation strategies for the time-varying VE function. We considered a Cox regression model estimated using partial likelihood for the vaccine-preventable outcome, which ignored the vaccine-irrelevant infections. The Cox model used a correctly-specified (log) linear model for the time-varying VE function $f(\cdot)$, however, the standard Cox model was misspecified because it did not account for the frailties. The other two estimators were the semiparametric logistic regression (SLR) model using the sieve estimation and TMLE estimation frameworks, with correctly specified linear models for $f(\cdot)$. The sieve estimator approximated the unknown nuisance parameter $\alpha_0(t)$ using a cubic B-spline model with number of knots equal to the cubic root of the number of events. The TMLE estimator estimated $\alpha_0(t)$ and each of the $r_0(T)$ functions using generalized additive model with thin-plate splines with default knots choices and degrees of freedom. After each TMLE update step, estimates of the linear predictors and $\alpha_0(t)$ were updated using the clever covariate using a profile approach which first updated the linear predictor and

then updated $\alpha_0(T)$. After each update of $\alpha_0(t)$, we re-smoothed $\alpha_0(t)$ in calendar time using a generalized additive model.

The results of our numerical experiments can be found in Tables B.1 and B.2. In Table B.1, we summarize the empirical coverage of 95% confidence intervals across 4 different bias scenarios – low heterogeneity in infection risk, high heterogeneity in infection risk, behavioral disinhibition, and vaccinating vulnerable groups early. Confidence interval coverage for Cox and SLR estimators were both comparable in settings with low risk heterogeneity and were close to the nominal 95% level. However, in the setting with high risk heterogeneity, the Cox model had coverage estimates that were far below the nominal level while both SLR models had estimated coverages very close to the nominal level. In the remaining two scenarios – behavioral disinhibition and a hidden association between risk and early vaccination, the standard Cox model had near zero coverage probability while the SLR models maintained 95% coverage. In Table B.2, we observe that all estimators had comparable mean squared errors in the low risk heterogeneity setting, but that the Cox model showed larger bias in the presence of high risk heterogeneity, behavioral disinhibition, and non-random vaccination dates.

Table B.1: Empirical 95% confidence interval coverage probabilities across 1000 replicates of different hypothetical COVID-19 vaccine studies. In all trials, $\lambda_{01} = 0.06$ and boost probability is $p = 0.85$ drawn from uniform distribution $[0, 1]$. Each simulated study enrolled $n = 10{,}000$.

				Constant VE			**Linear VE**		
Vaccination Assignment	Disinhibition	Force of Vaccine-Preventable infection (λ_{01})	Risk Heterogeneity (σ_U)	Cox	SLR (Sieve)	SLR (TMLE)	Cox	SLR (Sieve)	SLR (TMLE)
Random	N	0.06	Low ($\sigma_U = 1$)	93.8	96.9	96.9	93.8	95.3	95.4
Random	N	0.06	High ($\sigma_U = 2$)	38.2	95.7	95.7	53.7	94.0	94.1
Random	Y	0.06	Low ($\sigma_U = 1$)	0.4	96.8	96.8	0.5	96.1	96.1
Vulnerable	N	0.06	Low ($\sigma_U = 1$)	0.5	96.2	96.3	0.4	95.0	94.9

Additionally, we sought to understand the impact of force of vaccine-preventable infection on bias in the time-varying VE curve. The results are shown in Table B.3. When there was low heterogeneity in infection risk, higher force of vaccine-preventable infections slightly eroded the confidence interval coverage for the standard Cox model, while SLR models exhibited robust coverage. When risk heterogeneity was high, confidence intervals for the standard Cox model were far below the nominal level and rapidly eroded with higher force of infection. In every case, the SLR models

Table B.2: Empirical mean squared error across 1000 replicates of different hypothetical COVID-19 vaccine studies. In all trials, $\lambda_{01} = 0.06$ and boost probability is $p = 0.85$ drawn from uniform distribution $[0, 1]$. Each simulated study enrolled $n = 10{,}000$.

Vaccination Assignment	Disinhibition	Force of Vaccine-Preventable infection (λ_{01})	Risk Heterogeneity (σ_U)	**Constant VE** Cox	SLR (Sieve)	SLR (TMLE)	**Linear VE** Cox	SLR (Sieve)	SLR (TMLE)
Random	N	0.06	Low ($\sigma_U = 1$)	0.040	0.042	0.041	0.034	0.050	0.049
Random	N	0.06	High ($\sigma_U = 2$)	0.097	0.032	0.032	0.044	0.037	0.036
Random	Y	0.06	Low ($\sigma_U = 1$)	1.041	0.024	0.023	0.988	0.026	0.026
Vulnerable	N	0.06	Low ($\sigma_U = 1$)	2.821	0.043	0.042	2.911	0.048	0.047

exhibited robust coverage near the nominal level.

Table B.3: Empirical coverage probabilities when $\lambda_{01} = 0.06$ and boost probability is $p = 0.85$ uniformly over $[0, 1]$. $n = 10{,}000$.

Vaccination Assignment	Disinhibition	Force of Vaccine-Preventable infection (λ_{01})	Risk Heterogeneity (σ_U)	Constant VE Cox	SLR (Sieve)	SLR (TMLE)	Linear VE Cox	SLR (Sieve)	SLR (TMLE)
Random	N	0.03	Low ($\sigma_U = 1$)	94.8	96.1	96.2	94.4	94.7	94.8
Random	N	0.06	Low ($\sigma_U = 1$)	93.8	96.9	96.9	93.8	95.3	95.4
Random	N	0.12	Low ($\sigma_U = 1$)	89.4	96.4	96.4	90.3	94.8	94.8
Random	N	0.03	High ($\sigma_U = 2$)	66.1	96.6	96.7	75.5	95.3	95.4
Random	N	0.06	High ($\sigma_U = 2$)	38.2	95.7	95.7	53.7	94.0	94.1
Random	N	0.12	High ($\sigma_U = 2$)	22.9	96.1	96.0	28.6	95.2	95.1

B.8 Extensions

B.8.1 Ensuring monotonicity of time-varying VE

Domain knowledge supports that after a short period of time for the vaccine to produce an immune response, vaccine efficacy can either remain constant or decrease, but cannot increase in time-since-vaccination. Hence, after a brief period, we expect $f(\cdot)$ to be monotonically non-decreasing. We would like our estimator to respect this constraint to ensure sensible interpretation, mitigate finite-sample bias, and improve precision.

Assume $f(\tau) = \beta_0^T \underline{\psi}(\tau)$ for some known basis expansion $\{\psi_j\}_{j=1}^d$. Define $\hat{f}(\tau) = \hat{\beta}^T \underline{\psi}(\tau)$ as the unconstrained estimate of time-varying VE, where $\hat{\beta}$ is estimated from sieve/TMLE methods. Let $\mathcal{T} := \{\tau_1, \ldots, \tau_L\}$ denote a grid of L time-since-vaccination values of interest. Let $\hat{f}_l(\tau)$ and $\hat{f}_u(\tau)$

denote the upper and lower 95% confidence intervals of the unconstrained estimator over $\mathcal{T}$ which can be computed using the multivariate delta method.

Following results by Westling et al.[10], we can obtain a corrected monotonized estimator, $\hat{f}_{\text{mono}}(\tau)$, by simply running isotonic regression of the unconstrained estimator $\hat{f}$ over the grid $\mathcal{T}$. This is equivalent to projecting the unconstrained estimator onto the space of estimators satisfying $f(\tau_1) \leq f(\tau_2)$ for any $\tau_1 \leq \tau_2$. We can apply the same isotonic regression procedure to the unconstrained estimator's confidence limits ($\hat{f}_l(\tau)$, $\hat{f}_u(\tau)$) to obtain valid confidence limits for the monotonized estimator. In practice, we use precision-weighted isotonic regression to ensure that noisy values of $\hat{f}$ do not dominate the smoothing step. Westling et al.[10] established that the estimation error of $\hat{f}_{\text{mono}}$ over the grid is never worse than $\hat{f}$ and that the monotonized confidence intervals have no worse coverage than the original confidence intervals, and comes at no cost to the bandwidth over the grid.

B.8.2 Overcoming positivity violations by incorporating external disease data

In some cases, there may exist low identifiability for time-varying VE. This may occur in instances where it is hard to distinguish between time-varying vaccine effectiveness and changing compositions of infection types in the population over time. To illustrate, consider an extreme example is where every participant is vaccinated immediately upon enrollment. In this case, time-since-vaccination and calendar time are identical, and without additional information, we cannot distinguish between waning vaccine effectiveness and changing compositions of vaccine-preventable and vaccine-irrelevant infections over time; hence, time-varying VE is not identified. A more realistic example is a 1 year study where the vast majority of participants are enrolled and vaccinated within a few months. In this example, while time-varying VE is identified, it may be difficult to distinguish between time-varying vaccine effects and fluctuating incidence of infections in the broader population. In this situation, we say that time-varying VE may exhibit low identifiability.

The approaches we consider leverage external, population-level pathogen surveillance data to profile out the nuisance parameter describing the background relative fluctuations in disease incidence ($\alpha_0(t)$) and concentrate the available information on identifying the time-varying VE.

In our first case, we consider an analysis where the analyst has access to an external, population-

level disease surveillance data (hereafter denoted as 'ext') for both the vaccine-preventable ($J = 1$) and vaccine-irrelevant ($J = 0$) pathogens. Suppose that adopt the following *trend transport (TT1)* assumption, which claims for all $t \geq 0$,

$$\text{(TT1)} \qquad \alpha_0(t) = \log\left\{\frac{h_{01}(t)}{h_{00}(t)}\right\} = \log\left\{\frac{h^{\text{ext}}_{01}(t)}{h^{\text{ext}}_{00}(t)}\right\}$$

The TT1 assumption implies that the change-of-scale parameter which describes the time-varying relative differences in baseline hazard between the vaccine-preventable and vaccine-irrelevant infection types also applies to the external population. Importantly, TT1 allows the hazards of infection in the study population to be arbitrarily higher or lower than the external population, as long as the *baseline hazard ratio* for the two outcomes is the same. This is a critical point, since vaccine studies may be conducted in populations with intrinsically higher/lower risk of infection than the broader population.

Given that population-level case counts are usually summarized over discrete time intervals (per day or week), a natural estimator of the hazard in the external population on date/week t is the nonparametric maximum likelihood estimator of the discrete hazard function, $\hat{h}^{\text{ext}}_{0j}(t) = \frac{D^{\text{ext}}_j(t)}{R^{\text{ext}}_j(t)}$, where $D^{\text{ext}}_j(t), R^{\text{ext}}_j(t)$ are the number of incident outcomes of type j detected at date/week t and the number of people in the external population at risk of outcome j on date/week t. In most population-level contexts, the size of risk set is very large and minimally impacted by incident cases. Making the very mild assumption that the size of the risk sets for COVID-19 and off-target infection are of similar magnitudes and are much larger than the number of incident cases, we can write the change-of-scale parameter as follows.

$$\alpha_0(t) \overset{\text{(TT1)}}{=} \log\left(\frac{h^{\text{ext}}_{01}(t)}{h^{\text{ext}}_{00}(t)}\right) \approx \log\left(\frac{D^{\text{ext}}_1(t)/R^{\text{ext}}_1(t)}{D^{\text{ext}}_0(t)/R^{\text{ext}}_0(t)}\right) \approx \log\left(\frac{D^{\text{ext}}_1(t)}{D^{\text{ext}}_0(t)}\right)$$

The unknown nuisance parameter can be well-approximated by the (log) ratio of incident vaccine-preventable to vaccine-irrelevant cases at time t. Hence, $\log\left(D^{\text{ext}}_1(t)/D^{\text{ext}}_0(t)\right)$ can be substituted in for $\alpha_0(t)$ in a model for the odds of infection, thereby eliminating the nuisance parameter from the model. Hence, even in an entirely vaccinated cohort, the time-varying VE function can still be identified under the TT1 assumption.

In our second case, we consider an analysis where population-level disease surveillance data exists for the vaccine-preventable ($J = 1$) but **not** the vaccine-irrelevant ($J = 0$) pathogen. This may be

the case where the vaccine-preventable pathogen is of special public-health interest and vaccine-irrelevant pathogens are not tracked. In this case, may be reasonable to assume that the incidence of vaccine-preventable pathogen varies much more in calendar time than the vaccine-irrelevant pathogen, which may exhibit more stable incidence. We formalize this intuition by making the *stable disease assumption (SDA)* on vaccine-irrelevant infection and an alternative *trend transport (TT2)* assumption for the vaccine-preventable infection. Formally,

$$\text{(SDA)} \qquad h_{00}(t) = a_0 \; \forall t \geq 0 \qquad\qquad \text{(TT2)} \qquad \frac{h_{01}(t_1)}{h_{01}(t_2)} = \frac{h_{01}^{\text{ext}}(t_1)}{h_{01}^{\text{ext}}(t_2)} \; \forall t_1, t_2 \geq 0$$

SDA assumes that the hazard for off-target endpoint is constant over calendar time, although this can be relaxed to allow predictable (seasonal) variations. TT2 allows the hazard for COVID-19 to differ between study and external populations as long as the *calendar-time trend* in COVID-19 incidence is the same between the two populations. Define $\alpha_0 := \log\left(\frac{\left[\prod_{i=1}^{M} h_{01}^{\text{ext}}(t_i)\right]^{1/M}}{h_{00}(t)}\right)$. We can express the change-of-scale parameter as

$$\begin{aligned}
\alpha(t) &\overset{\text{(SDA)}}{=} \log\left(\frac{h_{01}(t)}{\left[\prod_{i=1}^{M} h_{01}(t_i)\right]^{1/M}}\right) + \alpha_0 \\
&\overset{\text{(TT2)}}{=} \log\left(\frac{h_{01}^{\text{ext}}(t)}{\left[\prod_{i=1}^{M} h_{01}^{\text{ext}}(t_i)\right]^{1/M}}\right) + \alpha_0 \\
&\approx \log\left(\frac{D_1^{\text{ext}}(t)}{\left[\prod_{i=1}^{M} D_1^{\text{ext}}(t_i)\right]^{1/M}}\right) + \alpha_0
\end{aligned}$$

Where M indexes a collection of calendar dates, $\{t_1, \ldots, t_M\}$ where surveillance data were available. Hence, the variation in $\alpha(t)$ over calendar time depends only on the trend in incident cases of the vaccine-preventable pathogen in the external population. The SDA and TT2 assumptions simplify the task of estimating a function-valued nuisance parameter $\alpha_0(t)$ to a single scalar α_0, which significantly boosts identifiability of $f(\cdot)$ in the logistic regression model.

B.8.3 Strain-specific time-varying VE

Suppose the vaccine-preventable outcome consists of $m \geq 2$ distinct subtypes, such as infections caused by different strains or variants of the vaccine-preventable infection. We may wish to estimate

time-varying VE against each variant separately. One strategy is to adopt a competing risk framework to accommodate different strain-specific hazard functions[11]. However, this approach would require estimating a unique change-of-scale parameter for each strain relative to the off-target infection, $\alpha_s(t) := \log\{h_{01s}(t)/h_{00}(t)\}$ where $s \in \{1, \ldots, m\}$. Moreover, the change-of-scale parameters are only well-defined over time periods where a strain s and the vaccine-irrelevant pathogen co-circulate, which may be violated during periods where older variants disappear and new variants emerge. Our approach takes inspiration from the work by Larson and Dinse[12] by modeling the unconditional (on infection) distribution of vaccine-preventable ($J = 1$) versus vaccine-irrelevant ($J = 0$) infections. The we model the conditional (on vaccine-preventable infection $J = 1$) distribution of failure types caused by different strains.

Let S refer to a categorical random variable denoting the strain of vaccine-preventable infection taking values in $\{1, \ldots, m\}$. For simplicity, we assume the following the multiplicative frailty models hold.

$$\begin{aligned} h_0(t) &= h_{00}(t)U(t) \\ h_{1s}(t) &= h_{01s}(t)U(t)\exp\{I(V \leq t)[f_s(t-V)]\} \end{aligned}$$

Suppose we define $h_{01}(t) := \sum_{k=1}^{m} h_{01k}(t)$ as the sum of each of the baseline hazards over the m vaccine-preventable infection subtypes. By multiplying and dividing by $h_{01}(t)$ and applying the change-of-scale parameter $\alpha_0(t) := \log\{h_{01}(t)/h_{00}(t)\}$ governing the relationship between any vaccine-preventable infection and off-target infection, we can rewrite the hazards above as follows.

$$\begin{aligned} h_0(t) &= h_{00}(t)U(t) \\ h_{1s}(t) &= h_{00}(t)p_{0s}(t)U(t)\exp\{I(V \leq t)f_s(t-V) + \alpha_0(t)\} \end{aligned}$$

Where $p_{0s}(t) := \frac{h_{01s}(t)}{\sum_{k=1}^{m} h_{01k}(t)}$ is the softmax probability of experiencing an vaccine-preventable infection with strain s conditional on experiencing a vaccine-preventable infection at time t and not being vaccinated $Z(t) = 0$ and $\log\{U(t)\} = 0$.

Another way to define $p_{0s}(t)$ is by noting that conditional on unvaccinated participant i experiencing a vaccine-preventable infection $J_i = 1$, the probability that the infection is of strain $S = k$

is given by

$$\begin{aligned}\frac{h_{1k}(t)}{\sum_{s'=1}^{m} h_{1s'}(t)} &= \frac{h_{00}(t)p_{0k}(t)U(t)\exp\{\alpha_0(t)\}}{\sum_{s'=1}^{m} h_{00}(t)p_{0s'}(t)U(t)\exp\{\alpha_0(t)\}} \\ &= \frac{\cancel{h_{00}(t)U(t)\exp\{\alpha_0(t)\}} \cdot p_{0k}(t)}{\cancel{h_{00}(t)U(t)\exp\{\alpha_0(t)\}} \cdot \sum_{s'=1}^{m} p_{0s'}(t)} \\ &= \frac{p_{0k}(t)}{\sum_{s'}^{m} p_{0s'}(t)} = p_{0k}(t)\end{aligned}$$

Where the first equality follows from plugging in the hazard parametrizations, the second follows from grouping like terms, the third equality follows from canceling the like terms, and the fourth equality follows from recognizing that $\sum_{s'}^{m} p_{0s'}(t) = 1$. Hence, $p_{0s}(t)$ can be interpreted as $P(S = s|T = t, J = 1, Z = 0)$, the probability of failure from strain $S = s$ among vaccine-preventable infections ($J = 1$) occurring at time $T = t$ among unvaccinated ($Z(t) = 0$) participants. As before, we can reparametrize the vaccine-preventable hazards according to $\alpha_0(t) = \log\{h_{01}(t)/h_{00}(t)\}$ which refers to the baseline hazard ratio between *any* vaccine-preventable infection and vaccine-irrelevant infection. The variant mixture proportions $p_s(t)$ parametrize the differences between the strain-specific vaccine-preventable hazards.

Given there are $S \in \{1, \ldots, m\}$ possible outcomes, we could imagine running m different logistic regressions comparing vaccine-preventable infections of strain s to vaccine-irrelevant infections.

$$\log\left\{\frac{P(J = 1, S = s|V = v, T = t)}{P(J = 0, S = 0|V = v, T = t)}\right\} = I(v \leq t)f_s(t - v) + \alpha_0(t) + \log\{p_s(t)\}$$

Where $\log(p_s(t))$ is an offset. Under this model, conditional on infection at time $T = t$, the probability of vaccine-preventable infection of strain s is equal to the softmax probability output.

$$P(J = 1, S = s|V = v, T = t) = \frac{p_s(t)\exp\{I(v \leq t)f_s(t - v) + \alpha_0(t)\}}{1 + \sum_{s'=1}^{m} p_{s'}(t)\exp\{I(v \leq t)f_{s'}(t - v) + \alpha_0(t)\}} \tag{B.1}$$

We briefly describe a few strategies on how to estimate the strain-specific mixture probabilities $\{p_s(t)\}_{s=1}^{m}$. One strategy is to leverage population-level disease surveillance data to obtain estimates of the strain-specific mixture probabilities, an approach referred to as surveillance anchored sieve analysis by Follmann et al.[13]. For example, NextStrain[14] or GISAID[15] provide high-resolution, population-level data on proportions of breakthrough infections caused by different viral variants. Since these are massive, population-level data sources, we can assume these variant probabilities are estimated with negligible error. As highlighted by Follmann et al.[13], external data may be

especially useful in identifying model parameters when the majority of participants are vaccinated, since it is impossible to distinguish between waning efficacy against a particular variant and the reduction in the variant's prevalence in the population. If the variant-specific probabilities from the external data sources $p_s^{\text{ext}}(t)$ are believed to be good approximations of the study-specific variant mixture proportions among unvaccinated infections $p_s(t)$, the externally-derived probabilities can be plugged in for the unknown $p_s(t)$ and treated as fixed. This is the strategy we will pursue in the COVE case study.

In principle, it is also possible to treat $\{p_s(t)\}_{s=1}^m$ as random and conduct estimation and inference estimate using a joint likelihood/estimating equation approach. This is left as an area of future research.

In the remaining subsections, we briefly describe how to conduct inference under the sieve estimation and TMLE approaches under multiple strains of vaccine-preventable infections when the strain-specific mixture probabilities $\{p_s(t)\}_{s=1}^m$ are known.

Sieve estimation under semiparametric multinomial logistic regression model

As before, suppose we approximate the change-of-scale parameter using a basis expansion that grows with the sample size.

$$\alpha_0(t) \approx \sum_{j=1}^{M} \alpha_j \phi_{j,v}(t)$$

We assume variant mixture proportions, $\{p_s(t)\}_{s=1}^m$, are borrowed from a population-level surveillance database and are assumed fixed. Moreover, suppose that the time-varying VE function against strain s is given by $f_s(\tau) = \beta_{0s}^T \underline{\psi}_s(t-v)$ where $\underline{\psi}$ is a d-dimensional basis and β_{0s} is the Euclidean parameter of interest. Hence, our model represents a semiparametric multinomial logistic regression model.

We rewrite the nonparametric multinomial likelihood in B.1 according to our semiparametric model as follows

$$P(J=0 \mid V=v, T=t) = \frac{1}{1+\sum_{s'=1}^m p_{s'}(t) \exp\{I(v \le t)\beta_{0s'}^T \underline{\psi}_{s'}(t-v) + \sum_{j=1}^M \alpha_j \phi_{j,v}(t)\}}$$

$$P(J=1, S=s \mid V=v, T=t) = \frac{p_s(t) \exp\{I(v \le t)\beta_{0s}^T \underline{\psi}_s(t-v) + \sum_{j=1}^M \alpha_j \phi_{j,v}(t)\}}{1+\sum_{s'=1}^m p_{s'}(t) \exp\{I(v \le t)\beta_{0s'}^T \underline{\psi}_{s'}(t-v) + \sum_{j=1}^M \alpha_j \phi_{j,v}(t)\}}$$

Then estimation and inference on $(\beta_{0s}^T\}_{s=1}^m$ and α_0^T can be accomplished by maximizing the multinomial likelihood associated with the models above.

TMLE under semiparametric multinomial logistic regression model

We provide a brief overview of the approach. Recall that we can define our semiparametric model for our data as follows.

$$P(J=1,S=s|V=v,T=t)=\frac{p_s(t)\exp\{I(v\leq t)f_s(t-v)+\alpha_0(t)\}}{1+\sum_{s'=1}^m p_{s'}(t)\exp\{I(v\leq t)f_{s'}(t-v)+\alpha_0(t)\}}$$
$$P(J=0|V=v,T=t)=\frac{1}{1+\sum_{s'=1}^m p_{s'}(t)\exp\{I(v\leq t)f_{s'}(t-v)+\alpha_0(t)\}}$$

Let $Y_s=I(J=1,S=s)$ and let $Y_0=I(J=0)$. Define the vector $\underline{Y}=(Y_1,\ldots,Y_m)^T$ and define the corresponding unconditional event probabilities based on the formula above.

$$Q_{0s}(t,v):=P(J=1,S=s|V=v,T=t)$$
$$Q_0(t,v):=\sum_{s=1}^m Q_{0s}(t,v)$$

The multinomial logistic regression model is equivalent to a series of binary regressions, meaning for $s=1,\ldots,m$.

$$\log\left\{\frac{Q_{0s}(t,v)}{1-Q_0(t,v)}\right\}=\log\{p_s(t)\}+I(v\leq t)f_s(t-v_s)+\alpha_0(t)$$

Where $\alpha_0(t)$ is an unknown nuisance, $p_s(t)$ are known offsets, and $f_s(\cdot\_s)$ is known the time-varying VE function against strain s which is known up to a scalar parameter β_s.

Suppose that $f_s(t-v_s)=\underline{\psi}_s(t-v)\beta_s$. The log-likelihood of a single contribution is given by.

$$\ell=\sum_{s=1}^m Y_s\{\log\{p_s(T)\}+I(V\leq T)\underline{\psi}_s(T-V)\beta_s+\alpha_0(t)\}$$
$$-\log\left\{1+\sum_{r=1}^m p_r(T)\exp\{I(V\leq T)\underline{\psi}_r(T-V)\beta_r+\alpha_0(t)\}\right\}$$

The *parametric* score with respect to β_s is given by

$$S_{\beta_s}(O)=I(V\leq T)\underline{\psi}_s(T-V)\left(Y_s-Q_{0s}(T,V)\right)$$

which follows from multinomial calculus.

To identify the efficient influence curve, we note that an influence function that is orthogonal to the nuisance tangent space (i.e., scores obtained by fluctuation α_0) is the unique EIF. Next up is to derive the nuisance tangent space of $\alpha_0(t)$. The nuisance tangent space will take the following form.

$$T_{nuis,\alpha_0}(P_0) = \{h(T)\sum_{r=1}^{m}(Y_r - Q_{0r}(V,T)) : h \in L_2\}$$

Note that we wish to obtain the *efficient* score by taking the parametric score and subtracting off its projection onto the nuisance tangent space. Note that the projection of $S_{\beta_s,0}$ onto T_{nuis,α_0} will take the form

$$a_s(T)\sum_{r=1}^{m}(Y_r - Q_{0r}(V,T))$$

Where the particular choice of $a(t)$ is such that the residual is orthogonal to every $h(t)$. That is,

$$\begin{aligned}
a_s(t) &= \frac{\mathbb{E}[S_{\beta_{0s}}\sum_{r=1}^{m}(Y_r - Q_{0r})|T=t]}{\mathbb{E}[(\sum_{r=1}^{m}(Y_r - Q_{0r}))^2\,|T=t]} \\
&= \frac{\mathbb{E}[I(V\leq T)\underline{\psi}_s(T-V)(Y_s - Q_{0s})\sum_{r=1}^{m}(Y_r - Q_{0r})|T=t]}{\mathbb{E}[(\sum_{r=1}^{m}(Y_r - Q_{0r}))^2\,|T=t]} \\
&= \frac{\mathbb{E}[I(V\leq T)\underline{\psi}_s(T-V)(Y_s - Q_{0s})(J - Q_0)|T=t]}{\mathbb{E}[(J-Q_0)^2\,|T=t]} \\
&= \frac{\mathbb{E}[I(V\leq T)\underline{\psi}_s(T-V)\mathbb{E}[(Y_s - Q_{0s})(J-Q_0)\mid|T=t,V]\mid T=t]}{\mathbb{E}[\mathbb{E}[(J-Q_0)^2\mid T=t,V]\mid T=t]} \\
&= \frac{\mathbb{E}[I(V\leq T)\underline{\psi}_s(T-V)\mathbb{E}[Y_s - Y_sQ_0 - JQ_{0s} + Q_{0s}Q_0\mid|T=t,V]\mid T=t]}{\mathbb{E}[\mathbb{E}[J-2JQ_0+Q_0^2\mid T=t,V]\mid T=t]} \\
&= \frac{\mathbb{E}_V[I(V\leq t)\underline{\psi}_s(t-V)Q_{0s}(V,t)(1-Q_0(V,t))\mid T=t]}{\mathbb{E}_V[Q_0(V,t)(1-Q_0(V,t))\mid T=t]}
\end{aligned}$$

Where the first equality is to achieve orthogonality, the second follows from plugging in the formula for the score, the third follows from noting that $J = \sum_{r=1}^{m} Y_r$ and $Q_0 = \sum_{r=1}^{m} Q_{0r}$, the fourth from iterated expectation, the fifth follows from expanding the products and simplifying binary terms, and the sixth follows from evaluating the inner expectations.

Hence, the efficient score function for β_{0s} is given by

$$\begin{aligned}
S_{\text{eff},\beta_{0s}}(O) &= I(V\leq T)\underline{\psi}_s(T-V)\{Y_s - Q_{0s}(T,V)\} - a_s(T)\sum_{r=1}^{m}(Y_r - Q_{0r}(V,T)) \\
&= I(V\leq T)\underline{\psi}_s(T-V)\{Y_s - Q_{0s}(T,V)\} - a_s(T)(J - Q_0(V,T))
\end{aligned}$$

The EIC is proportional to the efficient score up to a scaling matrix. The EIC can be derived by scaling the efficient score by the inverse of its variance.

Before we derive the TMLE algorithm, it is important to note that a multinomial logistic regression model can be expressed as a series of independent binary regressions which compare each vaccine-preventable infection with strain S to the "pivot" outcome of vaccine-irrelevant infections. As above, we can write the model as

$$\log\left(\frac{Q_{0s}(t,v)}{1-Q_0(t,v)}\right) = \log\{p_s(t)\} + I(v \leq t) f_s(t - v_s) + \alpha_0(t)$$

The TMLE algorithm associated with our model is writeable as follows.

(C1) Produce initial estimators of $\{\hat{Q}_{0s}^{(0)}\}_{s=1}^m$ by estimating $\alpha_0(t)$ and $\{\beta_{0s}\}_{s=1}^m$, using initial estimators $\hat{\alpha}_0(t)$ and $\hat{\beta}_s^{(0)}$, and treating the variant mixture proportions as known.

(C2) Estimate the unknown nuisance parameters $\{a_s(T)\}_{s=1}^m$. We propose estimating this quantity using a Nadaraya–Watson (NW) local kernel smoothing estimator, which essentially assigns a local (kernel) weight to values of T close to t. For instance, the N-W estimator of $a_S(T)$ is given by

$$\hat{a}_s^{NW}(t) = \frac{\frac{1}{n}\sum_{i=1}^n K_h(T_i - t) I(V_i \leq t) \underline{\psi}_s(t - V_i) \hat{Q}_{0s}^{(0)}(V_i, t)(1 - \hat{Q}_0^{(0)}(V_i, t))}{\frac{1}{n}\sum_{i=1}^n K_h(T_i - t) \hat{Q}_0^{(0)}(V_i, t)(1 - \hat{Q}_0^{(0)}(V_i, t))}$$

Where $K_h(x) = \frac{1}{h} K\left(\frac{x}{h}\right)$ and K is the kernel shape. For instance, a Gaussian kernel $K(x) = \frac{1}{\sqrt{2\pi}} \exp\{-x^2/2\}$.

(C3) For each s, construct the clever covariate vector.

$$H_s(T, V) := I(V \leq T) \underline{\psi}_s(T - V) - \hat{a}_s^{NW}(T)$$

(C4) Using the available study data, fluctuate the initial model predictions using a multinomial logit model which includes the original model predictions as offsets.

$$\log\left(\frac{Q_{0s}^{\epsilon}(t,v)}{1-Q_0^{\epsilon}(t,v)}\right) = \log\left(\frac{\hat{Q}_{0s}^{(0)}(t,v)}{1-\hat{Q}_0^{(0)}(t,v)}\right) + \epsilon_s H_s$$

(C5) For each $s \in \{1, \ldots, m\}$, update the parameter estimates using the coefficient on the as-

sociated clever covariates: $\hat{\beta}_s^{(1)} = \hat{\beta}_s^{(0)} + \hat{\epsilon}_s H_s$. Define the updated parameter estimates $\hat{\beta}^{(1)} := (\hat{\beta}_{s=1}^{(1)}, \ldots, \hat{\beta}_{s=m}^{(1)})$.

(C6) Using and $\hat{\alpha}_0(t)$ and $\hat{\beta}^{(1)}$, compute the updated event probabilities $\hat{Q}_{0s}^{(1)}$ using the multinomial model above.

(C7) Iterate until the updates $(\epsilon_1, \ldots, \epsilon_m)$ are sufficiently small (below some tolerance level), which implies that the mean of the EIC is close to zero.

Bibliography

[1] JD Holt and RL Prentice. Survival analyses in twin studies and matched pair experiments. *Biometrika*, 61, 1974.

[2] S. A. Murphy and A. W. Van Der Vaart. On Profile Likelihood. *Journal of the American Statistical Association*, 95(450):449–465, June 2000. ISSN 0162-1459, 1537-274X. doi: 10.1080/01621459.2000.10474219. URL `http://www.tandfonline.com/doi/abs/10.1080/01621459.2000.10474219`.

[3] Nan Xuan Lin and William Edward Henley. Prior event rate ratio adjustment for hidden confounding in observational studies of treatment effectiveness: a pairwise cox likelihood approach. *Statistics in Medicine*, 39:5149–5169, 2016.

[4] Jerald F. Lawless. *Statistical Models and Methods for Lifetime Data*. Jon Wiley and Sons Inc, 2003.

[5] Søren Johansen. An Extension of Cox's Regression Model. *International Statistical Review / Revue Internationale de Statistique*, 51(2):165, August 1983. ISSN 03067734. doi: 10.2307/1402746. URL `https://www.jstor.org/stable/1402746?origin=crossref`.

[6] Hung Chen. Asymptotically Efficient Estimation in Semiparametric Generalized Linear Models. *The Annals of Statistics*, 23(4):1102–1129, August 1995. ISSN 0090-5364, 2168-8966. doi: 10.1214/aos/1176324700. URL `https://projecteuclid.org/journals/annals-of-statistics/volume-23/issue-4/`

Asymptotically-Efficient-Estimation-in-Semiparametric-Generalized-Linear-Models/10.1214/aos/1176324700.full.

[7] Lars van der Laan and Peter Gilbert. Semiparametric logistic regression for inference on relative vaccine efficacy in case-only studies with informative missingness. arXiv preprint arXiv:2303.11462, 2025. URL https://arxiv.org/abs/2303.11462.

[8] Mark J van der Laan and Sherri Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York, 2011.

[9] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21, 2018.

[10] Ted Westling, Mark J. van der Laan, and Marco Carone. Correcting an estimator of a multivariate monotone function with isotonic regression. *Electronic Journal of Statistics*, 14(2):3032–3069, January 2020. ISSN 1935-7524, 1935-7524. doi: 10.1214/20-EJS1740. URL https://projecteuclid.org/journals/electronic-journal-of-statistics/volume-14/issue-2/Correcting-an-estimator-of-a-multivariate-monotone-function-with-isotonic/10.1214/20-EJS1740.full.

[11] RL Prentice, JD Kalbfleisch, AV Peterson Jr, N Flournoy, et al. The analysis of failure times in the presence of competing risks. *Biometrics*, 34(4):541–554, December 1978.

[12] Martin G Larson and Gregg E Dinse. A mixture model for the regression analysis of competing risks data. *J.R. Stat. Soc. Ser. C. Appl. Stat.*, 34(3):201–211, 1985.

[13] Dean Follmann, Michael Fay, and Craig Magaret. Estimation of vaccine efficacy for variants that emerge after the placebo group is vaccinated. *Statistics in Medicine*, 41(16):3076–3089, July 2022.

[14] James Hadfield, Colin Megill, Sidney M Bell, John Huddleston, Barney Potter, Charlton Callender, Pavel Sagulenko, Trevor Bedford, and Richard A Neher. Nextstrain: real-time track-

ing of pathogen evolution. *Bioinformatics*, 34(23):4121–4123, December 2018. ISSN 1367-4803, 1367-4811. doi: 10.1093/bioinformatics/bty407. URL https://academic.oup.com/bioinformatics/article/34/23/4121/5001388.

[15] Shruti Khare, Céline Gurry, Lucas Freitas, Mark B Schultz, et al. GISAID's Role in Pandemic Response. *China CDC Weekly*, 3(49):1049–1051, 2021. ISSN 2096-7071. doi: 10.46234/ccdcw2021.255. URL http://weekly.chinacdc.cn/en/article/doi/10.46234/ccdcw2021.255.

Appendix C

SUPPLEMENTARY MATERIALS (CHAPTER 4)

C.1 Additional Simulation Results

C.1.1 Performance Under Violations of NCO Assumption

As described in the main text of the paper, we explored the performance of NCO adjustment when the NCO assumption (Assumption 4) was violated to varying degrees. We assessed performance in terms of mean absolute bias relative to the plug-in estimator and coverage of nominal confidence intervals as a function of the magnitude of the NCO violation, sample size, and the predictiveness of the NCO. We varied the effect of treatment on the candidate NCO within $\beta_N \in \{0, 0.1, 0.25, 0.5, 0.75, 1, 1.5, 2\}$, encompassing cases where the NCO assumption is not violated and is violated to varying degrees. For simplicity, we focused on comparing the plug-in estimator to the NCO-adjusted estimator. We considered pairing the NCO-adjusted estimator with either a pretest of the sharp null hypothesis or an equivalence pretest at various equivalence thresholds $\epsilon > 0$. All pretests were conducted at level $\alpha = 0.05$. If the sharp null hypothesis was rejected, the plug-in estimator would be returned, otherwise the NCO-adjusted estimator was returned. If the equivalence null hypothesis was rejected, the NCO-adjusted estimator was returned and otherwise, the plug-in estimator was returned. The effect on the primary outcome was set to $\beta = 1$, and $\rho_{Y,X} = 0.3$ and $\rho_{(Y,N|X)}$ was varied between 0.3 and 0.8. We focused on the identity link function as in Setting 1. Estimators were compared on the basis of bias and confidence interval coverage.

Results for the simulation are shown in Figure C.1. Adjustment for the candidate NCO can lead to substantial relative bias as the effect of treatment on the candidate NCO increases. Imposing sharp null/equivalence pretests acts to mitigate the bias when the NCO assumption fails. However, most pretests are vulnerable to intermediate violations of the NCO assumption for small sample sizes. Under large violations of the NCO assumption, the pretest flags the violation and avoids bias. Under small violations, NCO adjustment proceeds as before but generates minimal bias. The performance of equivalence tests depends on the choice of equivalence threshold ϵ. Small choices of equivalence threshold set very stringent thresholds to permit NCO adjustment, and therefore

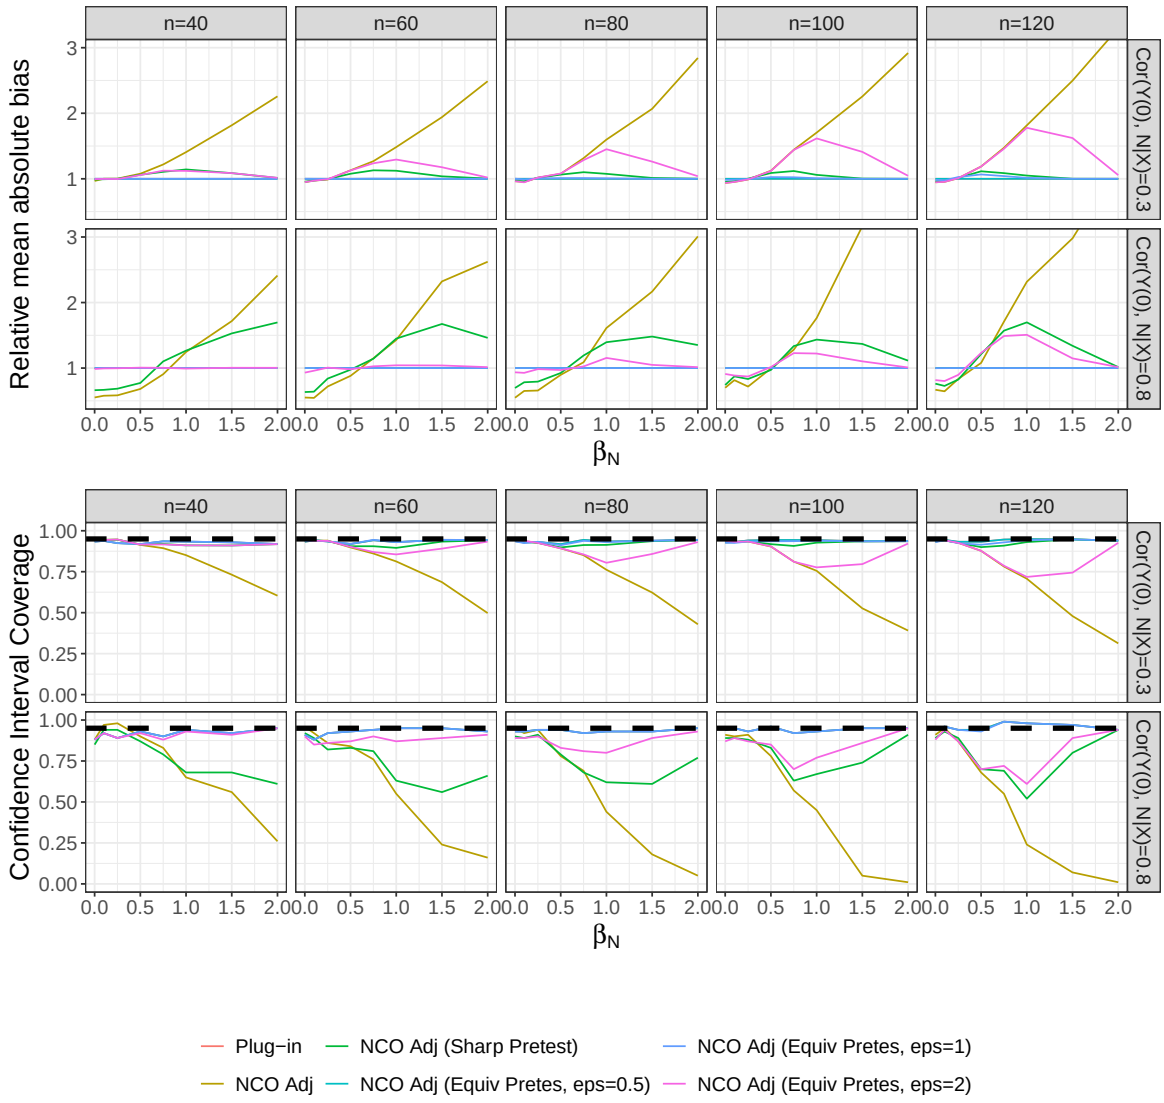

Figure C.1: Numerical experiments illustrating the impact of NCO violations on the performance of NCO adjustment. Performance summarized using relative mean absolute bias and confidence interval coverage. Key parameters varied are sample size, correlation between $Y(0)$ and N conditional on X, and the effect of treatment on the candidate NCO β_N.

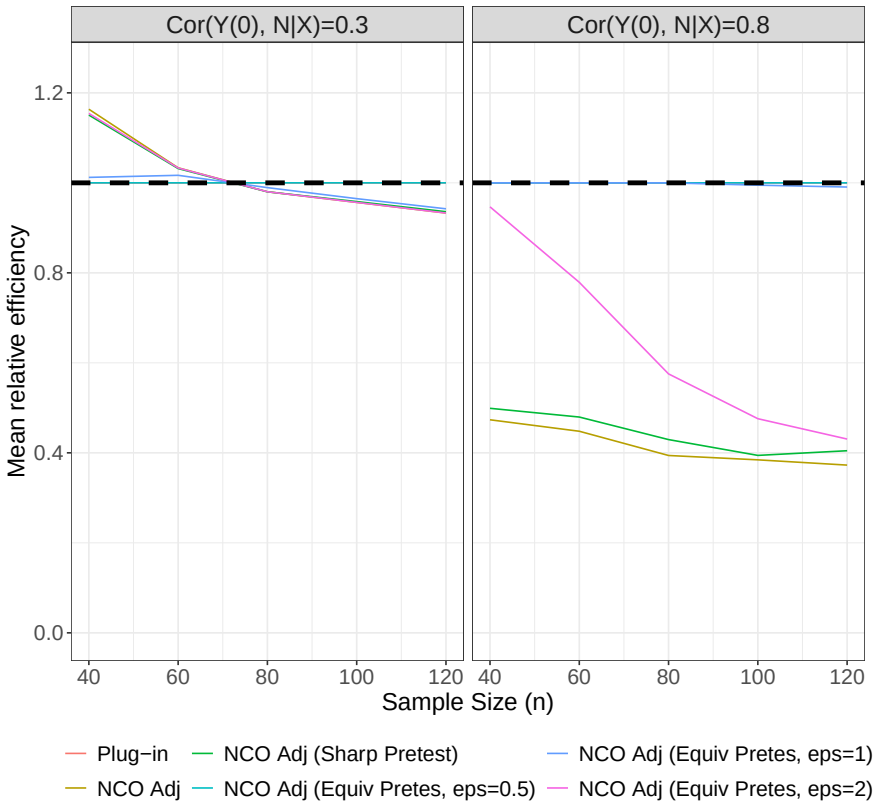

Figure C.2: Performance of NCO adjustment (with and without pretests) when the NCO assumption is satisfied ($\beta_N = 0$). Key parameters varies are sample size (n) and correlation between $Y(0)$ and N conditional on X.

rarely adjust for the NCO and result in an estimator which mimics the plug-in. Larger equivalence thresholds only protect against violations of the NCO assumption that exceed the threshold and are more vulnerable to intermediate violations of the NCO assumption. The results for confidence interval coverage mirror those for the bias. NCO adjustment provides valid coverage when the NCO assumption is satisifed, but has very low coverage when the treatment has greater effect on the NCO candidate. Incorporating pretests guards against cases where NCO adjustment may be very harmful. As before, the performance of equivalence pretests depend on whether the violation of the NCO lies within or beyond the equivalence threshold.

Next, we compare the performance of the different estimators when the NCO assumption is satisifed to see whether there is any penalty paid for pretests in Figure C.2. In the case where the

NCO is not very predictive, adjustment for the NCO leads to loss of precision in finite samples and a small precision gain in larger studies. When the NCO is highly prognostic, NCO adjustment leads to substantial gains in precision. In general, pretests reduce the benefit of NCO adjustment. Tests of the sharp causal null lead to a small precision loss relative to NCO adjustment outright, because the default assumption made by the sharp null is that the candidate NCO is valid. Most equivalence tests led to substantial precision losses, save for the equivalence test with the widest equivalence margin. As sample size increased, this equivalence pretest favored more frequent NCO adjustment, and average efficiency relative to the plug-in estimator decreased.

C.1.2 Results for Inference on SATE and Randomization Tests

In Figure C.3, we illustrate the performance of OLS-adjusted estimators of the sample average treatment effect (SATE) in numerical experiments in settings consistent with the main simulation described in the main text of the article. Similar to estimators of the population average treatment effect, we see that adjustment for the NCO can reduce finite-sample bias, increase precision, and increase power of Wald tests especially when the NCO becomes more prognostic for the primary outcome. When the NCO distribution is skewed, adjustment for the NCO on the quantile scale can mitigate the bias, efficiency loss, and power loss of estimators which adjust for the raw NCO. We observe that adjustment of the standard errors using the HC3 correction leads to improved coverage of confidence intervals in finite samples. Power to reject the weak null hypothesis $H_0^{\text{weak}} : \text{SATE} = 0$ using robust t-statistics us shown in the bottom right panel for $n = 60$. We observe that NCO-adjustment can lead to improvements in power to detect non-zero treatment effects when the NCO is more prognostic for the primary outcome.

In Figure C.4, we illustrate the power of unadjusted and adjusted randomization tests of $H_0^{\text{sharp}} : Y(1) = Y(0)$ across simulated datasets with sample size $n = 60$ and randomization ratio $\pi = 0.80$. We see that when N and Y exhibit very low correlations, the adjusted and unadjusted approaches have very similar power. However, when N and Y exhibit moderate to high correlation, adjustment for N can lead to increases in statistical power to test the sharp causal null. In settings with the skewed predictor subject to detection limits, adjusting for the raw NCO on the quantile scale led to reduced power relative to covariate-adjusted and plug-in approaches. However, adjustment for the

NCO on the quantile scale led to superior statistical power among all estimators. All estimators produced nominal Type I error rates under the null hypothesis of no treatment effect.

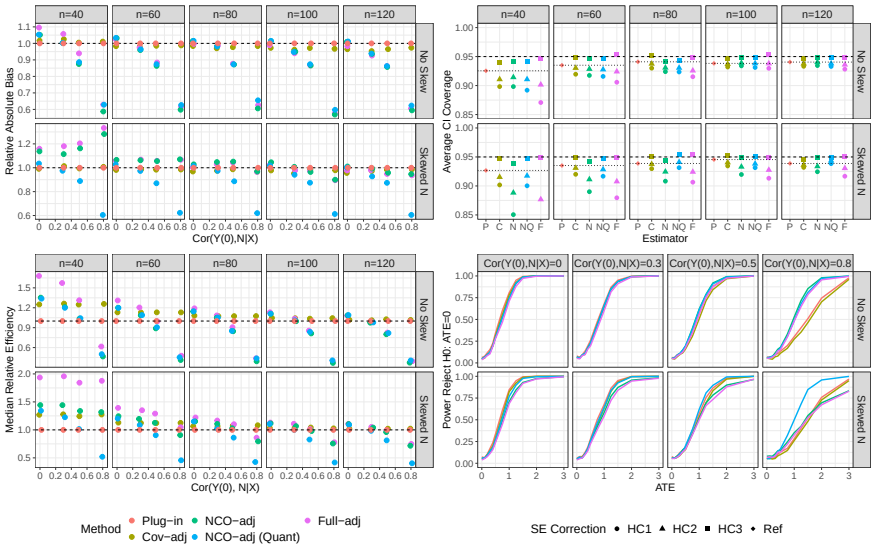

Figure C.3: Top left: relative mean absolute bias of different estimators of the SATE as a function of sample size, data generating mechanism, and how prognostic the NCO is for the outcome of interest. Top right: average coverage of 95% confidence intervals of different estimators of the SATE and finite sample standard error corrections as a function of sample size and data generating mechanism. Bottom left: median relative efficiency of estimators of the SATE as a function of sample size, data generating mechanism, and how prognostic the NCO is for the outcome of interest. Bottom right: power curves of Wald tests of H0: SATE=0 as a function of how prognostic the NCO is for the outcome of interest and the data generating mechanism.

C.1.3 Identities linking manipulable parameters to parameters of data generating process

Recalling the data generating process in Section 4.3 of the main text. Note that

$$\begin{pmatrix} Y \\ N \end{pmatrix} \Big| X \sim x\beta_1 + N(U\beta_2, \Sigma = \text{diag}(1,2))$$

$$U|X \sim N(0,1)$$

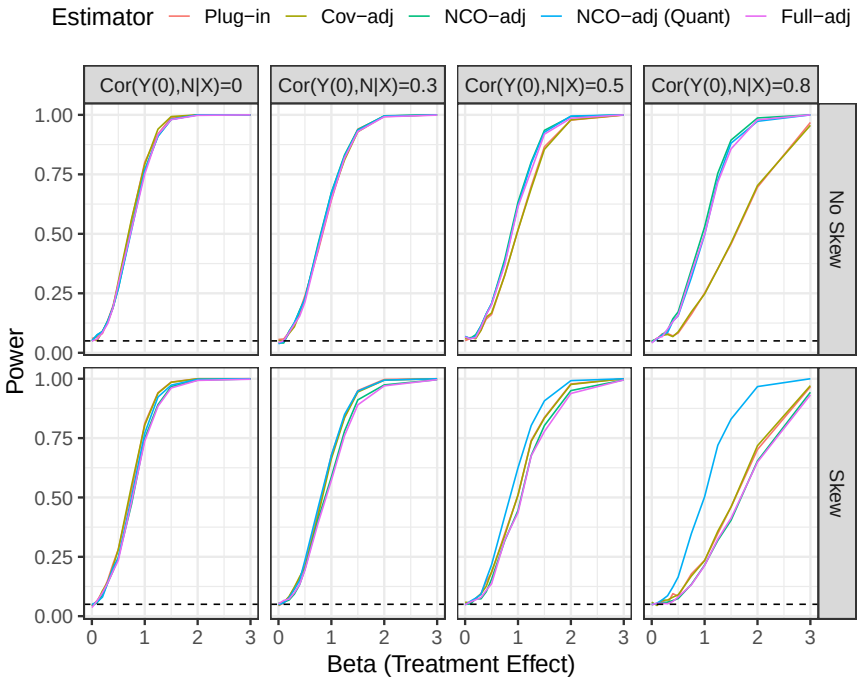

Figure C.4: Power of randomization tests of $H_0^{\text{sharp}} : Y(1) = Y(0)$ across simulated datasets as a function of residual correlation between Y and N, working model specification, and treatment effect.

A key parameter that we want to modulate between simulations is the partial correlation between Y and N conditional on X, defined as

$$\rho_{Y,N|X} = \text{Cov}(Y, N|X)/\sqrt{\text{Var}(Y|X)\text{Var}(\text{N—X})}$$

Using laws of total variance and covariance, we can express each of these unknowns according to the regression parameters

$$\text{Var}(Y|X) = \text{Var}(E[Y|X,U]|X) + E[\text{Var}(Y|X,U)|X] = \text{Var}(X\beta_1 + U\beta_2|X) + 1 = \beta_2^2 + 1$$

$$\text{Var}(N|X) = \text{Var}(E[N|X,U]|X) + E[\text{Var}(N|X,U)|X] = \text{Var}(X\beta_1 + U\beta_2|X) + 1 = \beta_2^2 + 1$$

$$\text{Cov}(Y, N|X) = \text{Cov}(E(Y|U,X), E(Y|U,X)\ |X) + \text{E}[\text{Cov}(Y, N|U,X)|X] = \text{Var}(X\beta_1 + U\beta_2\ |X) + 0 = \beta_2^2$$

Hence,

$$\implies \rho_{Y,N|X} = \frac{\beta_2^2}{1+\beta_2^2}$$
$$\beta_2 = \sqrt{\frac{\rho_{Y,N|X}}{1-\rho_{Y,N|X}}}$$

Next, we examine

$$(X,Y)^T \sim N\left((0, X\beta_1 + U\beta_2), \Sigma = \begin{pmatrix} 1 & \rho_{X,Y} \\ \rho_{X,Y} & 1 \end{pmatrix}\right)$$
$$X \sim N(0,1)$$

The unknown parameter can be expressed as

$$\rho_{X,Y} = \mathrm{Cov}(X,Y)/\sqrt{\mathrm{Var}(Y)\mathrm{Var(X)}}$$

Note that we know that $\mathrm{Var}(X) = 1$. Now we can solve for the remaining components using laws of total variance/covariance.

$$\mathrm{Var}(Y) = \mathrm{Var}(E[Y|X,U]) + \mathrm{E}[\mathrm{Var}(Y|X,U)] = \mathrm{Var}(X\beta_1 + U\beta_2) + 1 = \beta_1^2 + \beta_2^2 + 1$$
$$\mathrm{Cov}(X,Y) = \mathrm{Cov}(E[X|X,U], E[Y|X,U]) + E[\mathrm{Cov}(X,Y|X,U)] = \mathrm{Cov}(X, \beta_1 X + \beta_2 U) + 0 = \beta_1$$

Solving we obtain

$$\rho_{X,Y} = \beta_1/\sqrt{\beta_1^2 + \beta_2^2 + 1}$$
$$\implies \beta_1^2 = \rho_{X,Y}^2(\beta_1^2 + \beta_2^2 + 1)$$
$$\implies \beta_1 = \sqrt{\frac{\rho_{X,Y}^2(\beta_2^2+1)}{1-\rho_{X,Y}^2}}$$

C.2 Proofs of propositions

C.2.1 Background on semiparametric theory

We provide a short primer on semiparametric theory which defines the useful terminology and concepts. In many classical statistical problems, we are interested in estimating a target quantity within a statistical model M restricted to distributions indexed by a finite dimensional parameter

θ. Such models are referred to as *parametric models*. However, in many cases, we may not want to rely on a possibly-misspecified parametric model for inference for fear of incurring systematic bias. Instead, we we may restrict attention to distributions indexed by infinite dimensional parameters. A model M indexed by infinite-dimensional parameter without any assumptions is referred to as the *nonparametric model*. However, there are multitude of other restrictions that we could place on our model M which may be motivated by scientific or experimental knowledge. A common set of assumptions we make, particularly in causal inference problems, are exclusion restrictions, or assumptions of independence between variables. For example, in a randomized trial measuring baseline covariates X, we can restrict focus to data generating laws satisfying the exclusion restriction $X \perp A$ which is satisfied by ignorability/randomization. We refer to such models with some arbitrary restriction on the infinite-dimensional parameter as *semiparametric models*.

In parametric models, maximum likelihood estimators are celebrated due to their *efficiency* in large samples, meaning they asymptotically attain the Cramer-Rao lower bound, or the minimum possible variance an estimator can achieve under the parametric model. In semiparametric estimation problems, we are concerned with finding analogs to the parametric case. First, we are interested in what is the minimum possible variance that we can estimate a a target quantity across all data generating laws in the model. We refer to this quantity as the *semiparametric efficiency bound*. Second, we are interested in identifying what estimator which asymptotically achieves this minimum variance bound. We refer to this quantity as the *semiparametric efficient estimator*.

We will restrict focus to regular asymptotically linear (RAL) estimators to avoid pathological estimators which display very unstable behavior uniformly over the parameter set. Parametric efficiency theory can be generalized to nonparametric and semiparametric statistical models by relying on the insight that the best possible variance attainable in an infinite-dimensional model M should be *at least as large* as the best possible variance attainable in *any* parametric submodel $M_1 \subset M$. To obtain the tightest bound, we claim that the best possible variance in an infinite dimensional model is equal to the smallest achievable variance in the least favorable quadratic mean differentiable parametric submodel. Let $v_0^*(M)$ denote the variance of any RAL estimator in an infinite dimensional model. Let $\mathcal{H}(P_0)$ denote the collection of all smooth (quadratic mean differentiable) parametric submodels centered at P_0. We restrict focus to quadratic mean differentiable models as they will have score functions with mean zero and bounded variance. Let M_h denote a particular

choice of submodel. The generalized Cramer-Rao lower bound is as follows.

$$v_0^*(M) \geq \sup_{h \in \mathcal{H}(P_0)} v_0(M_h) = \sup_{h \in \mathcal{H}(P_0)} \frac{\left(\frac{\partial}{\partial \theta} \psi(P_{\theta,h}) \Big|_{\theta=0} \right)^2}{P_0 g_h^2}$$

where g_h is the score of the least-favorable parametric model through h. $P_0 g_h^2$ is the variance of the score function in the least favorable submodel. The numerator is a square of the pathwise derivative of the functional at P_0 evaluated along the submodel h. If ψ is a pathwise differentiable functional, Riesz representation theorem guarantees the pathwise derivative is writable as an inner product of a gradient function $D(P_0)$ in the Hilbert space $L_0^2(P_0)$ and the score g_h along the least favorable submodel. This yields a new and useful formulation of the efficiency bound.

$$v_0 * (M) \geq \sup_{h \in \mathcal{H}(P_0)} \frac{(P_0[D(P_0) g_h])^2}{P_0 g_h^2}$$

Recognizing that the variance bound depends on the set of parametric submodels $\mathcal{H}$ entirely through the score function g_h, we can rewrite the above bound in terms of the set of allowable scores $\mathcal{G}(P_0) := \{g_h : h \in \mathcal{H}(P_0)\}$, also known as the tangent space.

$$v_0^*(M) \geq \sup_{g \in \mathcal{G}(P_0)} \frac{(P_0[D(P_0) g])^2}{P_0 g^2}$$

The above bound will have a closed form when the gradient $D(P_0) \in \mathcal{G}(P_0)$, meaning the gradient lies in the tangent space. In a nonparametric model, the set of allowable scores is all of $\mathcal{G}(P_0) = L_0^2(P_0)$, therefore, $D(P_0) \in \mathcal{G}(P_0)$ by default and the nonparametric efficiency bound is given by

$$v_0^*(M) \geq P_0[D(P_0)^2]$$

However, when we impose restrictions to our model M, we thin the collection of allowable scores in $\mathcal{G}(P_0)$ and the nonparametric gradient $D(P_0)$ is not guaranteed to lie in $\mathcal{G}(P_0)$ anymore. However, we can represent any gradient in the following manner.

$$D(P_0) := D^*(P_0) + (D(P_0) - D^*(P_0))$$

Where $D^*(P_0) \in \mathcal{G}(P_0)$ and $D(P_0) - D^*(P_0) \in \mathcal{G}(P_0)^\perp$ represents the orthogonal complement of the tangent space. We refer to $D^*(P_0)$ as the *canonical gradient*. By substituting this expression

into the generalized Cramer-Rao bound above

$$v_0 * (M) \geq \sup_{g \in \mathcal{G}(P_0)} \frac{(P_0[(D^*(P_0) + (D(P_0) - D^*(P_0)))g])^2}{P_0 g^2}$$
$$= P_0[D^*(P_0)^2]$$

Which holds because $P_0((D(P_0) - D^*(P_0))g) = 0$ because $D(P_0) - D^*(P_0) \in \mathcal{G}(P_0)^\perp$, meaning it is orthogonal to any score function $g \in \mathcal{G}(P_0)$.

This is a critical insight: the generalized Cramer-Rao lower bound in a semiparametric model depends on the *orthogonal projection* of the gradient $D(P_0)$ onto the tangent space $\mathcal{G}(P_0)$. Hence, the recipe for obtaining the semiparametric efficiency bound is (a) identify a gradient in the nonparametric model $L^2(P_0)$ (which is guaranteed to be the unique gradient trivially by the uniqueness of a projection), (b) derive the form of the tangent space of allowable score functions in a semiparameteric model under some restrictions, and (c) compute the canonical gradient – or the projection of the nonparametric gradient onto the tangent space comprised of scores compatible with the semiparametric model. The variance of the projection will represent the semiparametric efficiency bound.

Once these steps are completed, obtaining the efficient estimator in the semiparametric model is straightforward. A key insight in semiparametric estimation theory is that there is 1-1 correspondence between influence functions of RAL estimators and gradients of pathwise differentiable parameters. Hence, constructing a RAL estimator with influence function equal to the canonical gradient will produce the semiparametric efficient estimator of the target quantity.

C.2.2 Proof of Proposition 3

The following theorem uses the results from the previous subsection to derive the form of the efficiency bound and efficient influence function (EIF).

Proof: Let A denote a binary treatment. Let X denote a covariate, N denote a valid NCO satisfying assumption 4, and Y denote the outcome. Suppose Assumption 1-3 hold. hence, the following exclusion restriction holds: $A \perp (X, N, Y(0), Y(1))$. We define our statistical model M as the set of all distributions for $(A, X, N, Y) \overset{iid}{\sim} P_0 \in M$ under the exclusion restriction above.

Suppose our target parameter is the ATE, which is identified under Assumptions 1-3.

$$\mathbb{E}[Y(1) - Y(0)] = \mathbb{E}[Y|A = 1] - \mathbb{E}[Y|A = 0]$$

To derive the efficiency bound, we follow the recipe described in the previous subsection. Since the ATE is pathwise differentiable, there exists a correspondence between influence functions of RAL estimators and gradients. Indeed, one can show that the influence function for the nonparametric/plug-in estimator is

$$D(P_0)(a, y) = \frac{a(y - \mathbb{E}[Y|A = 1])}{\pi} - \frac{(1 - a)(y - \mathbb{E}[Y|A = 0])}{1 - \pi}$$

We assume a known treatment probability π such that the sample estimate $\pi = \pi$. One can also substitute a consistent estimator π for π. The above influence function is also a gradient for our target parameter.

Step two in our recipe is to derive the form of the tangent space of M, i.e., the set of allowable score functions consistent with our model. We can decompose the into submodels for each constituent components.

$$M = M_{X,N} \otimes M_{A|X,N} \otimes M_{Y|X,N,A}$$

Note that under our exclusion restriction $A \perp (X, N)$, $M_{A|X,N} = M_A$. This model for A actually contains only a single density, $p(A|X, M) = \pi^A(1 - \pi)^{1-A}$, where π is known or consistently estimated. Since this component of the model consists of a single density, we can ignore this portion of the model when calculating the tangent space. We leave the other components of the model unspecified. Since we restricted focus to QMD models, we know the scores consist of all mean-zero, finite-variance functions. This implies they reside in the Hilbert space $L^2_0(P_0)$. Formally, the tangent space in model M can be written as an orthogonal sum

$$\mathcal{G} = \mathcal{G}_{X,N} \oplus \mathcal{G}_{Y|X,N,A}$$
$$\mathcal{G}_{X,N} := \{\text{all } g(X, N) \in L^2_{0,X,N}(P_0)\}$$
$$\mathcal{G}_{Y|X,N,A} := \{\text{all } g(X, N, Y, A) \in L^2_{Y|X,N,A}(P_0)\}$$

In the third step of our recipe, we must compute the projection of the gradient from step 1,

$D(P_0)$, onto the tangent space $\mathcal{G}$ to obtain the canonical gradient. Recall that $D(P_0) = D^*(P_0) + D^{\perp}(P_0)$. Hence, any gradient is writable as the canonical gradient plus an additional term that lies in the orthogonal complement of the tangent space. As pointed out by[1], any element of $\mathcal{G}$ can be written as

$$Ah_1(X,N) + (1-A)h_0(X,N)$$

for $h_1(X,N)$ and $h_0(X,N)$ both mean zero and finite variance. The moment restriction implied by randomization is

$$\begin{aligned}
&\mathbb{E}[Ah_1(X,N) + (1-A)h_0(X,N)|X,N] = 0 \\
&\implies h_1(X,N) = -\frac{\pi}{(1-\pi)}h_0(X,N)
\end{aligned}$$

Substituting into the original formula for an element of $\mathcal{G}$, we get that any element of $\mathcal{G}$ is writable as

$$\begin{aligned}
&A\left(-\frac{\pi}{(1-\pi)}h_0(X,N)\right) + (1-A)h_0(X,N) \\
&\equiv -A\pi h_0(X,N) + (1-A)(1-\pi)h_0(X,N) \\
&\equiv (A-\pi)h_0(X,N)
\end{aligned}$$

Hence, for a specific h, we can obtain the efficient influence function.

$$\frac{a(y - \mathbb{E}[Y|A=1])}{\pi} - \frac{(1-a)(y-\mathbb{E}[Y|A=0])}{1-\pi} - (A-\pi)h(X,N)$$

The specific h we are looking for is the projection of the nonparametric gradient onto the orthogonal complement of the tangent space. Using some useful projection properties, we obtain that the projection takes the form

$$\Pi(D(P_0)|\mathcal{G}^{\perp}) = \frac{\mathbb{E}[Y|X,N,A=1] - \mathbb{E}[Y|A=1]}{\pi} - \frac{\mathbb{E}[Y|X,N,A=0] - \mathbb{E}[Y|A=0]}{1-\pi}$$

Hence, the canonical gradient is given by

$$\begin{aligned}
D^*(P_0) = &\frac{a(y - \mathbb{E}[Y|A=1])}{\pi} - \frac{(1-a)(y-\mathbb{E}[Y|A=0])}{1-\pi} \\
&- (A-\pi)\left(\frac{\mathbb{E}[Y|X,N,A=1] - \mathbb{E}[Y|A=1]}{\pi} - \frac{\mathbb{E}[Y|X,N,A=0] - \mathbb{E}[Y|A=0]}{1-\pi}\right)
\end{aligned}$$

$$= D(P_0) - (A - \pi)\left(\frac{\mathbb{E}[Y|X, N, A = 1]}{\pi} - \frac{\mathbb{E}[Y|X, N, A = 0]}{1 - \pi}\right)$$

The semiparametric efficiency bound is given by the variance of this gradient

$$v_0^*(M_V) \geq P_0[D^*(P_0)^2]$$

By the 1-1 correspondence between gradients and influence functions of RAL estimators for pathwise differentiable parameters, we have that the RAL estimator that achieves the semiparametric efficiency bound, or the semiparametric efficient estimator of the ATE in a randomized trial with a NCO, takes the form

$$\frac{1}{n}\sum_{i=1}^{n} D^*(P_0)(X_i, N_i, A_i, Y_i) = \bar{Y}(1) - \bar{Y}(0) - \frac{1}{n}\sum_{i=1}^{n}(A_i - \pi)\left(\frac{\mathbb{E}(Y|X, N, A = 1)}{\pi} - \frac{\mathbb{E}(Y|X, N, A = 0)}{1 - \pi}\right)$$

We argue that the estimator which adjusts for (X, N) offers guaranteed efficiency gain over estimators which adjust for X alone. We argue this point using the following result: let M_1 be a statistical model satisfying $A \perp (X, N)$ and M_2 be a statistical model satisfying $A \perp X$. Clearly, $M_1 \subseteq M_2$. By properties of gradients in nested models, every gradient in M_2 is also a gradient in M_1. Hence, the canonical gradient (i.e., efficient influence function) in M_2 is a gradient in M_1, but may not necessarily be the canonical gradient/efficient influence function. Hence, the variance lower bound achievable in M_1 is at most variance lower bound achievable in M_2. □

C.2.3 Proof of Corollary 1

Consider the oracle augmented data generating mechanism which includes measurement of the latent variable U in the model where $(X, U, N, Y(1), Y(0)) \perp A$.

$$(A_i, X_i, U_i, N_i, Y_i) \overset{i.i.d}{\sim} P_0 \in M^*$$

Following the same logic as in the Proof of Proposition 1, the canonical gradient for the ATE in this expanded model where (X, U, N) is measured is given by

$$\begin{aligned} D_1^*(P_0) = &\frac{a(y - \mathbb{E}[Y|A = 1])}{\pi} - \frac{(1 - a)(y - \mathbb{E}[Y|A = 0])}{1 - \pi} \\ &- (A - \pi)\left(\frac{\mathbb{E}[Y|X, U, N, A = 1] - \mathbb{E}[Y|A = 1]}{\pi} - \frac{\mathbb{E}[Y|X, U, N, A = 0] - \mathbb{E}[Y|A = 0]}{1 - \pi}\right) \end{aligned}$$

$$= D(P_0) - (A - \pi)\left(\frac{\mathbb{E}[Y|X, U, N, A = 1]}{\pi} - \frac{\mathbb{E}[Y|X, U, N, A = 0]}{1 - \pi}\right)$$

Where $D(P_0)$ is the gradient in the nonparametric (unrestricted model).

Suppose we make the additional assumption that Y is *mean independent* of X, U conditional on N, A, or formally,

$$\mathbb{E}[Y|X, U, N, A = a] = \mathbb{E}[Y|N, A = a]$$

This assumption is related to the idea of surrogacy, as N is taken as a surrogate for the effect of (X, U) on Y. The canonical gradient in the randomized trial satisfying $A \perp (X, U, N)$ under the mean independence assumption is given by

$$D_2^*(P_0) = D(P_0) - (A - \pi)\left(\frac{\mathbb{E}[Y|N, A = 1]}{\pi} - \frac{\mathbb{E}[Y|N, A = 0]}{1 - \pi}\right)$$

The semiparametric efficiency bound, or the minimal variance in the semiparametric model of laws satisfying (i) $A \perp (X, U, N)$ and (ii) mean independence of Y and X, U conditional on N, A is equal to the second moment of the canonical gradient, $P_0[D_2^*(P_0)^2]$. Recognizing the 1-1 correspondence between gradients and influence functions of RAL estimators for pathwise differentiable parameters, we have that the RAL estimator with the canonical gradient as its influence function achieves the semiparametric efficiency bound. The semiparametric efficient estimator in the augmented data model under mean independence is given by the following.

$$\frac{1}{n}\sum_{i=1}^{n} D_2^*(P_0)(N_i, A_i, Y_i) = \frac{1}{n}\sum_{i=1}^{n} \frac{A_i Y_i}{\pi} - \frac{(1 - A_i)Y_i}{1 - \pi} - (A_i - \pi)\left(\frac{\mathbb{E}(Y|N, A = 1)}{\pi} - \frac{\mathbb{E}(Y|N, A = 0)}{1 - \pi}\right)$$

Furthermore, suppose that h_0 and h_1 are substituted for the unknown regression functions in the formula for the influence function. By virtue of belonging in the class of functions in Equation 4.3 , the estimator will be consistent and asymptotically normal regardless of choices of the working models. Suppose we use linear working models fit using OLS for the full (oracle) data unit and a reduced data unit that does not measured (X, U). Let the unknown parameters below denote the limits in probability of the parameters estimates obtained by OLS (note that estimates need not converge to the true parameters, merely some fixed limit).

$$h_a(X, U, N) = \beta_0 + \beta_1 X + \beta_2 U + \beta_3 N$$

$$h_a^*(N) = \beta_0^* + \beta_3^* N$$

We rely on the following idea that if two variables are independent, they are also uncorrelated. By assumption, $Y \perp (U, X)|(N, A)$, implying that $\beta_1 = \beta_2 = 0$ in the first working model. It follows that the two working models specified above are asymptotically equivalent, as the probability limits of the parameter estimates satisfy $\beta_0^* = \beta_0$ and $\beta_3 = \beta_3^*$ because $\beta_1 = \beta_2 = 0$. Hence, OLS working models based on the oracle data unit $\{X, U, N\}$ are asymptotically equivalent to OLS working models with N as the sole predictor, implying that $\hat{\psi}_{\text{NCO-AIPW}}$ is asymptotically equivalent to $\hat{\psi}_{\text{Oracle-AIPW}}$, the AIPW OLS-assisted estimator which adjusts for the oracle data unit.

As review, when using OLS regression models, $\hat{\psi}_{\text{Cov-AIPW}}$ is guaranteed to be more efficient than $\hat{\psi}_{\text{Plug-in}}$.[2] Adjusting for more explanatory variables cannot harm large-sample efficiency. Hence, $\hat{\psi}_{\text{Oracle-AIPW}}$ is guaranteed to be at least as efficient as $\hat{\psi}_{\text{Cov-AIPW}}$. As described above, because the probability limits of the working models $h_a(X, U, N)$ are equivalent to the probability limits of the working models $h_a^*(N)$, $\hat{\psi}_{\text{NCO-AIPW}}$ is asymptotically equivalent to $\hat{\psi}_{\text{Oracle-AIPW}}$. This implies that $\hat{\psi}_{\text{NCO-AIPW}}$ offers guaranteed efficiency gain relative to $\hat{\psi}_{\text{Cov-AIPW}}$ and no efficiency loss relative to $\hat{\psi}_{\text{Oracle-AIPW}}$ respectively, even when the working models are misspecified. □

C.3 Sensitivity Analysis

In the main portion of this paper, we propose using randomization tests of the sharp null or an equivalence test to identify violations of the NCO assumption. Below, we describe a sensitivity analysis based on linear structural models which could be used to investigate the impacts of assumption violations on the ATE.

C.3.1 Sensitivity Analysis Using Linear Models

We can pursue a sensitivity analysis to the NCO assumption based on linear models as described by Rosenbaum[3]. The bias of an estimator which adjusts for a variable affected by treatment in a randomized experiment is the result of marginalizing over the factual distribution of N rather than the proper counterfactual distributions of $N(1), N(0)$:

$$\text{Bias} = \mathbb{E}_N[\mathbb{E}_0[Y(1)|N = n] - \mathbb{E}_0[Y(0)|N = n]] - \left[\mathbb{E}_{N(1)}[\mathbb{E}_0[Y(1)|N = n] - \mathbb{E}_{N(0)}[\mathbb{E}_0[Y(0)|N = n]]\right]$$

$$\equiv (\mathbb{E}_N - \mathbb{E}_{N(1)})\{\mathbb{E}_0[Y(1)|N=n]\} - (\mathbb{E}_N - \mathbb{E}_{N(0)})\{\mathbb{E}_0[Y(0)|N=n]\}$$

Where $\mathbb{E}_N[\cdot]$ refers to expectation taken over the observed distribution of N and $\mathbb{E}_{N(a)}[\cdot]$ refers to expectation taken over the distribution of counterfactuals $N(a)$. When Assumption 4 holds and $N = N(1) = N(0)$, the bias term is equal to zero exactly.

Let $P_{N(a)}$ denote the measure of the counterfactual post-baseline outcome and $P_N := \alpha P_{N(1)} + \beta P_{N(0)}$ denote the measure of the observed distribution of N, which is necessary a mixture of the two counterfactual distributions such that $\alpha, \beta \in [0,1]$ and $\alpha + \beta = 1$. We can rewrite the expectations in empirical process form.

$$\text{Bias} = \int \mathbb{E}_0[Y(1)|N=n] d((\alpha-1)P_{N(1)} + \beta P_{N(0)}) - \int \mathbb{E}_0[Y(0)|N=n] d((\beta-1)P_{N(0)} + \alpha P_{N(1)}).$$

Suppose we assume the following parallel linear structural nested models for the outcomes.

$$\mathbb{E}_0[Y(1)|N=n] = \alpha_1 + \gamma N$$
$$\mathbb{E}_0[Y(0)|N=n] = \alpha_0 + \gamma N$$

The bias can be expressed as

$$\begin{aligned}\text{Bias} &= \int \alpha_1 + \gamma N \; d((\alpha-1)F_{N(1)} + \beta F_{N(0)}) - \int \alpha_0 + \gamma N \; d((\beta-1)F_{N(0)} + \alpha F_{N(1)}) \\ &= \gamma \int N d(F_{N(0)} - F_{N(1)}) = -\gamma \left(\mathbb{E}_0[N(1) - N(0)]\right).\end{aligned}$$

The bias of the adjusted estimator is linear in the average treatment effect on the post-baseline outcome. One can estimate γ by fitting an OLS estimators with main effects for (A, N) and no interaction, and using the estimated regression coefficient on N. Then, sensitivity analysis can be performed by choosing plausible values for the effect of treatment on the auxiliary outcomes, $\mathbb{E}_0[N(1) - N(0)]$, and subtracting the bias from the estimated treatment effect.

Bibliography

[1] Marie Davidian, Anastasios A. Tsiatis, and Selene Leon. Semiparametric Estimation of Treatment Effect in a Pretest–Posttest Study with Missing Data. *Stat Sci*, 20(3):261–301, August 2005. ISSN 0883-4237. doi: 10.1214/088342305000000151. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2600547/`.

[2] Anastasios A. Tsiatis, Marie Davidian, Min Zhang, and Xiaomin Lu. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Stat Med*, 27(23):4658–4677, October 2008. ISSN 0277-6715. doi: 10.1002/sim.3113. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2562926/`.

[3] Paul R. Rosenbaum. The consequences of adjustment for a concomitant variable that has been affected by the treatment. *J R Stat Soc Ser B*, 147, 1984.

Appendix D

SUPPLEMENTARY MATERIALS (CHAPTER 5)

D.1 Identification (Proposition 4)

First, we show that identification of $S_{a_0}(t)$ fails when using a G-formula which adjusts for coarsened versions of N such as $\tilde{N}$ and Δ_N. Without loss of generality, we can focus on adjusting for $\tilde{N}$, knowing that identical arguments will follow for the observed event indicator.

Recall the definition of our target parameter, $S_{a_0}(t) := P(Y(a_0) \geq t_0)$, which under randomization and the NCO assumption, can be written using the G-formula with respect to the full-data unit where N is observed.

$$S_{a_0}(t) = \mathbb{E}_{X,N}\left[S_{0Y}(t \mid A = a_0, X, N)\right]$$

Consider the following estimator which marginalizes over the distribution of X and the observed NCI time $\tilde{N}$.

$$\theta = \mathbb{E}_{X,\tilde{N}}\left[P(Y > t \mid A = a_0, X, \tilde{N})\right]$$

By causal consistency and strong ignorability, we have that

$$\mu_{a_0}(t, X, \tilde{n}) = P(Y > t \mid A = a_0, X, \tilde{N} = \tilde{n}) = P(Y(a_0) > t \mid X, \tilde{N}(a) = \tilde{n})$$

Where $\tilde{N}(a) := \min(N, C_N(a))$. Hence,

$$\theta \equiv \mathbb{E}_{X,\tilde{N}}\left[\mu_{a_0}(t, X, \tilde{N})\right]$$

To illustrate why θ is biased for the target $S_{a_0}(t)$, we appeal to the ideas proposed by Rosenbaum (1984) regarding adjustment for a post-treatment variable. In brief, the literature of principal stratification claims that the potential outcome under a specific arm, such as $\tilde{N}(a)$, is an intrinsic, deterministic characteristic of a participant, meaning it behaves exactly like a baseline covariate.[2] Because $\tilde{N}(a)$ is essentially a baseline covariate, standardizing with respect to the counterfactual

distribution must identify the marginal survival curve of interest.

$$S_{a_0}(t) = \mathbb{E}_{X,\tilde{N}(a_0)}\left[P(Y(a) \geq t \mid A = a_0, X, \tilde{N}(a_0))\right] \equiv \mathbb{E}_{X,\tilde{N}(a_0)}\left[\mu_{a_0}(t, X, \tilde{N}(a_0))\right]$$

Now we can see plainly why θ is biased for the target parameter. The failure of identification results from the improper marginalization of μ_{a_0} over the *observed law* of $(X, \tilde{N})$ instead of the *counterfactual law* of $(X, \tilde{N}(a_0))$. The observed law is a contaminated mixture of the two potential outcome distributions.

$$P(\tilde{N} \mid X) = \pi_0(a_0 \mid X)P(\tilde{N}(a_0) \mid X) + \pi_0(1 - a_0 \mid X)P(\tilde{N}(1 - a_0) \mid X)$$

In general, $\mathbb{E}_{X,\tilde{N}}[\mu_{a_0}(t, X, \tilde{N})] \neq \mathbb{E}_{X,\tilde{N}(a_0)}[\mu_{a_0}(t, X, \tilde{N}(a_0)]$ when the distribution of $\tilde{N} \overset{d}{\neq} \tilde{N}(a_0)$[1]. Hence, the failure of identification results from the fact that our target parameter marginalizes over the counterfactual distribution $\tilde{N}(a_0)$ while the proposed parameter incorrectly marginalizes over the distribution of the observed variable $\tilde{N}$.

Hence, identification using observed, coarsened versions of N, such as $\tilde{N}$ or Δ_N fail to identify the target parameter without additional assumptions. We prove that the following inverse probability weighted complete case (IPW) formula does identify the target parameter.

$$S_{a_0}(t_0) = \mathbb{E}_0\left[\frac{\Delta_N}{G_0(\tilde{N} \mid A, X)}\frac{I(A = a_0)}{\pi_0(a_0 \mid X)}P_0(Y > t \mid A = a_0, X, \tilde{N})\right]$$

Where $\mathbb{E}_0[\cdot]$ refers to marginalizing over the distribution of $(A, X, \Delta_N, \tilde{N})$.

$$\begin{aligned}
&\mathbb{E}_0\left[\frac{\Delta_N}{G_0(\tilde{N} \mid A, X)}\frac{I(A = a_0)}{\pi_0(a_0 \mid X)}P_0(Y > t \mid A = a_0, X, \tilde{N})\right] \\
&= \mathbb{E}_{A,X,N,C}\left[\frac{I(N \leq C)}{G_0(N \mid A, X)}\frac{I(A = a_0)}{\pi_0(a_0 \mid X)}P_0(Y > t \mid A = a_0, X, N)\right] \\
&= \mathbb{E}_{X,N}\left[P_A(A = a_0 \mid X, N) \times \mathbb{E}_C\left\{\frac{I(N \leq C)}{G_0(N \mid A = a_0, X)\pi_0(a_0 \mid X)}P_0(Y > t \mid A = a_0, X, N) \mid X, N, A = a_0\right\}\right] \\
&= \mathbb{E}_{X,N}\left[\mathbb{E}_C\left\{\frac{I(N \leq C)}{G_0(N \mid A = a_0, X)}P_0(Y > t \mid A = a_0, X, N) \mid X, N, A = a_0\right\}\right] \\
&= \mathbb{E}_{X,N}\left[P_0(Y > t \mid A = a_0, X, N)\cancel{\left\{\frac{G_0(N \mid A = a_0, X)}{G_0(N \mid A = a_0, X)}\right\}}\right] \\
&= \mathbb{E}_{X,N}\left[S_{0Y}(t \mid a_0, X, N)\right] = S_{a_0}(t)
\end{aligned}$$

The first equality follows from recognizing $\tilde{N} = N$ when $\Delta_N = 1$ with probability 1. The second

equality follows from Tower Law and law of total expectation. The third equality follows from canceling the treatment probabilities. The fourth follows from evaluating the expectation of the indicator function in the numerator and canceling like terms. The final equality follows from expressing the survival function $S_a(t)$ in G-computation form and using the arguments developed in Westling et al. (2023).

D.1.1 Remark 1

In the remark below, we discuss two important consequences of independent censoring (Assumption 20).

The first exclusion restriction follows directly from $(C_Y, C_N) \perp (Y, N) \mid (A, X)$. To show the latter exclusion restriction, we break it into two cases.

$$\{C_N \perp N \mid (A, X, \Delta_Y, \tilde{Y})\} = \begin{cases} C_N \perp N \mid (A, X, Y, C_Y > Y) & \text{when } \Delta_Y = 1 \\ C_N \perp N \mid (A, X, C_Y, Y \geq C_Y) & \text{when } \Delta_Y = 0 \end{cases}$$

We examine each of these cases separately to verify that Assumption 4 implies the latter result.

When $\Delta_Y = 1$, the conditional probability of C_N and N is given by

$$P(C_N, N \mid A, X, Y = y, C_Y \geq y) = \frac{P(C_N, N, C_Y \geq y, Y = y \mid A, X)}{P(C_Y \geq y, Y = y \mid A, X)}$$
$$= \frac{P(N, Y = y \mid A, X)}{P(Y = y \mid A, X)} \cdot \frac{P(C_N, C_Y \geq y \mid A, X)}{P(C_Y \geq y \mid A, X)}$$

Where the first equality follows from Bayes rule and the second from applying Assumption 4 to the numerator and denominator. Since the joint probability factors into a term depending on N and another depending on C_N, proving $C \perp N \mid A, X, \tilde{Y}, \Delta_Y = 1$.

A very similar result can be shown for the case when $\Delta_Y = 0$:

$$P(C_N, N \mid A, X, C_Y = c_Y, Y > c_Y) = \frac{P(C_N, N, C_Y = y, Y > c_Y \mid A, X)}{P(C_Y = c_Y, Y > c_Y \mid A, X)}$$
$$= \frac{P(N, Y > c_Y \mid A, X)}{P(Y > c_Y \mid A, X)} \cdot \frac{P(C_N, C_Y = c_Y \mid A, X)}{P(C_Y = c_Y \mid A, X)}$$

Where as before, Bayes rule and Assumption 4 show that the conditional likelihood can be factored into a term depending on N and another depending on C_N, which proves $C \perp N \mid A, X, \tilde{Y}, \Delta_Y = 0$.

Taken together, we see Assumption 20 implies $C_N \perp N \mid (A, X, \tilde{Y}, \Delta_Y)$.

D.2 Proof of Theorem 1

Recall that our target parameter is the treatment-arm-specific survivor function of Y at some landmark time t_0, $S_{a_0}(t_0)$. As a reminder, we refer to the "full" data as observations where the primary outcome Y is right-censored (by C_Y) by the predictive, negative control infection time N is observed for all units.

$$O^{\text{full}} := (A, X, \tilde{Y}, \Delta_Y, N)$$

In practice, we only sample the "observed" data, for which N is also censored (by C_N).

$$O^{\text{obs}} := (A, X, \tilde{Y}, \Delta_Y, \tilde{N}, \Delta_N)$$

How shall we estimate $S_{a_0}(t_0)$ efficiently using the observed data? We will base our estimate on an estimating equation built using the observed-data efficient influence (EIF), commonly referred to as a one-step estimator[4].

To establish the observed-data EIF, we must build on the theory of full data influence functions. Let $\mathcal{H}^F$ denote the full-data Hilbert space, which consists of mean-zero measurable functions of O^{full} with finite variance equipped with covariance inner product. Let $\mathcal{H}$ denote the observed-data Hilbert space, which consists of mean-zero measurable functions of O^{obs} with finite variance equipped with covariance inner product. We focus on these Hilbert spaces because they represent plausible choices of estimating functions that can be used to construct estimating equations for $S_{a_0}(t_0)$. Note that influence functions of $S_{a_0}(t_0)$ in our full-data model lie in a subspace of $\mathcal{H}$ orthogonal to the nuisance tangent space. Hence, influence functions represent choices of estimating functions that are minimally influenced by perturbations of the nuisance parameters, which make them ideal choices for estimation.

A direct application of Theorem 1 of Westling et al. (2023) yields the EIF in the full-data model where N is observed for all participants.

$$\phi^F_{a_0,t_0}(o^{\text{full}}) = S_{0Y}(t_0 \mid a_0, x, n)\left[1 - \frac{I(a = a_0)}{\pi_0(a_0 \mid x)}\left\{\frac{I(\tilde{y} \le t_0, \delta_y = 1)}{S_{0Y}(\tilde{y} \mid a, x, n)G_{0Y}(\tilde{y} \mid a, x)} - \int_0^{\min(t_0,\tilde{y})} \frac{\Lambda_{0Y}(du \mid a, x, n)}{S_{0Y}(u \mid a, x, n)G_{0Y}(u \mid a, x)}\right\}\right]$$

Note that $\pi_0(a_0 \mid x)$ does not depend on N by Assumption 5, and $G_{0Y}(u \mid a, x)$ does not depend on N by Assumptions 4 and 5. However, S_{0Y} does depend on N, and the additional regressor is what drives the precision gain.

How do we proceed from the full-data model to the observed-data model? We make this transition by recognizing that the censoring of N defines a coarsening of the full data unit, O^{full}. Following Tsiatis (2006), we represent the observed data as $O^{\text{obs}} = \{\mathcal{C}_N, G_{\mathcal{C}_N}(O^{\text{full}})\}$ where the coarsening level $\mathcal{C}_N$ is

$$\mathcal{C}_N = \begin{cases} C_N & \text{if } C_N < N \\ \infty & \text{if } C_N \geq N \end{cases}$$

and $G_r(O^{\text{full}})$ is the corresponding many-to-one function of the full data satisfying

$$G_r(O^{\text{full}}) = \begin{cases} (A, X, \tilde{Y}, \Delta_Y, C_N = r, N > r) & \text{if } r < \infty \\ (A, X, \tilde{Y}, \Delta_Y, C_N \geq N, N) & \text{if } r = \infty \end{cases}$$

Following Tsiatis (2006), we can express the coarsening variable's hazard at level r as

$$\lambda_r(O^{\text{full}}) = P(\mathcal{C}_N = r \mid \mathcal{C}_N \geq r, O^{\text{full}}).$$

Using the equivalence $\{\mathcal{C}_N = r\} \equiv (C_N \geq r, N > r) \cup (N \leq C_N)$; the events overlap on $\{r < N \leq \mathcal{C}_n\}$. We can rewrite the hazard in terms of the censoring variable C_N.

$$\lambda_r(O^{\text{full}}) = P(C_N = r, N > r \mid (C_N \geq r, N > r) \cup (N \leq C_N), O^{\text{full}})$$

On the event $\{N < r\}$, the hazard is zero, so we can restrict attention to the event $\{N \geq r\}$, on which the union in the conditioning event reduces to $\{C_N \geq r\}$.

$$\lambda_r(O^{\text{full}}) = I(N \geq r)P(C_N = r \mid C_N \geq r, O^{\text{full}}) = I(N \geq r)\lambda_{C_N}(r \mid O^{\text{full}})$$

Where $\lambda_{C_N}(r \mid O^{\text{full}}$ is the hazard of C_N at level r conditional on the full data unit. Under coarsening at random (5.3), we have that the $\lambda_{C_N}(r \mid O^{\text{full}}) = \lambda_{C_N}(r \mid G_r(O^{\text{full}}))$, yielding

$$\lambda_r(O^{\text{full}}) = I(N \geq r)\lambda_{C_N}(r \mid A, X, \tilde{Y}, \Delta_Y)$$

Where $\lambda_{C_N}(r \mid A, X, \tilde{Y}, \Delta_Y)$ is the hazard of C_N at level r conditional on $(A, X, \tilde{Y}, \Delta_Y)$.

Define the conditional survival and cumulative hazard functions of C_N given $(A, X, \tilde{Y}, \Delta_Y)$ as

follows:

$$G_{0N}(t \mid a, x, \tilde{y}, \delta_y) := \exp\left\{\int_0^t \lambda_{C_N}(u \mid a, x, \tilde{y}, \delta_y)du\right\}$$
$$\Lambda_{C_N}(t \mid a, x, \tilde{y}, \delta_y) := \int_0^t \lambda_{C_N}(u \mid a, x, \tilde{y}, \delta_y)du$$

Let $N_C(t) := I(\tilde{N} \leq t, \Delta_N = 0\}$ denote the counting process for censoring and let $Y_R(t) := I(\tilde{N} \geq t)$ denote the at risk indicator. Under 5.3 implied by Assumption 4, the process

$$M_{C_N}(t) := N_C(t) - \int_0^t Y_R(u)d\Lambda_{C_N}(u \mid a, x, \tilde{y}, \delta_y)$$

is a martingale with respect to the filtration generated by the observed data on the time scale of C_N.

By Theorem 10.4 of Tsiatis (2006) (in its continuous-time form, equation 10.76) and using the fact that $\phi^F_{a_0,t_0}$ has finite variance, the projection of $\phi^F_{a_0,t_0}$ onto the observed-data Hilbert space is as follows.

$$\phi^{\text{obs}}_{a_0,t_0} = \frac{\Delta_N \phi^F_{a_0,t_0}(\tilde{Y}, \Delta_Y, A, X, \tilde{N})}{G_{0N}(\tilde{N} \mid A, X, \tilde{Y}, \Delta_Y)} + \int_0^\infty \frac{h^*(u, A, X, \tilde{Y}, \Delta_Y)}{G_{0N}(u \mid A, X, \tilde{Y}, \Delta_Y)} dM_{C_N}(u \mid A, X, \tilde{Y}, \Delta_Y)$$

Where h^* is an augmentation function chosen to make $\phi^{\text{obs}}_{a_0,t_0}$ orthogonal to the nuisance tangent space contributed by the C_N coarsening mechanism. The optimal choice (in the sense of the observed data influence function with lowest variance) is

$$h^*(u, a, x, \tilde{y}, \delta_y) = Q_0(u \mid a, x, \tilde{y}, \delta_y) := \mathbb{E}[\phi^F_{a_0,t_0}(\tilde{Y}, \Delta_Y, A, X, N) \mid N \geq u, A = a, X = x, \tilde{Y} = \tilde{y}, \Delta_Y = \delta_y]$$

the conditional expectation of the full-data influence function over the at-risk set $\{N \geq u\}$, given baseline factors and primary outcome history.

Substituting Q_0 in for h^* and expanding the censoring martingale into two terms, $dM_{C_N}(u) = dN_C(u) - Y_R(u)d\Lambda_{C_N}(u)$, the integral term can be writen as

$$\int_0^\infty \frac{Q_0(u)}{G_{0N}(u)} dM_{C_N}(u) = \frac{(1 - \Delta_N)Q_0(\tilde{N})}{G_{0N}(\tilde{N} \mid A, X, \tilde{Y}, \Delta_Y)} - \int_0^{\tilde{N}} \frac{Q_0(u \mid A, X, \tilde{Y}, \Delta_Y)}{G_{0N}(u \mid A, X, \tilde{Y}, \Delta_Y)} d\Lambda_{C_N}(u \mid A, X, \tilde{Y}, \Delta_Y)$$

Where we express some conditioning for the sake of compactness. Combining with the result from Theorem 10.4 above,

$$\phi^{\text{obs}}_{a_0,t_0} = \frac{\Delta_N \phi^F_{a_0,t_0}(\tilde{Y}, \Delta_Y, A, X, \tilde{N})}{G_{0N}(\tilde{N} \mid A, X, \tilde{Y}, \Delta_Y)} + \frac{(1 - \Delta_N)Q_0(\tilde{N} \mid A, X, \tilde{Y}, \Delta_Y)}{G_{0N}(\tilde{N} \mid A, X, \tilde{Y}, \Delta_Y)}$$

$$-\int_0^{\tilde{N}} \frac{Q_0(u \mid A, X, \tilde{Y}, \Delta_Y)}{G_{0N}(u \mid A, X, \tilde{Y}, \Delta_Y)} d\Lambda_{C_N}(u \mid A, X, \tilde{Y}, \Delta_Y)$$

The expression has the form of an augmented inverse probability of censoring complete-case (AIPWCC) estimating function, matching the influence function showed in 5.4: the first term is the inverse-probability-weighted complete-case contribution from participants whose N is observed; the second and third terms together form the augmentation that recovers efficiency by exploiting partial information from N-censored participants.

Now we wish to show that the influence function $\phi^{\text{obs}}_{a_0,t_0}$ is the efficient influence function globally. Note that Theorem 10.4 supports that $\phi^{\text{obs}}_{a_0,t_0}$ is the "optimal" choice of observed-data influence function based on a particular choice of full-data influence function $\phi^F_{a_0,t_0}(\cdot)$. However, in general, there is no guarantee that our choice of $\phi^F_{a_0,t_0}(\cdot)$ leads to the optimal observed data influence function.

However, in this case, we can rightly claim that $\phi^{\text{obs}}_{a_0,t_0}$ is globally efficient. By Theorem 1 of Westling et al. (2023), $\phi^F_{t_0,a_0}$ is the *unique/nonparametric* EIF in the full-data model. Since the full-data influence function is unique, the projection onto the observed data Hilbert space is also unique. Hence, $\phi^{\text{obs}}_{a_0,t_0}$ must be the globally efficient influence function in the observed-data model. □

D.3 Proof of Theorem 2

Let $\eta^* = (\pi^*, G^*_Y, G^*_N, S^*_Y, f^*_N)$ denote the specified nuisance functions. Since $\phi^{\text{obs}}_{a_0,t_0}$ is written as an uncentered influence function, consistency follows if

$$\mathbb{E}\{\phi^{\text{obs}}_{a_0,t_0}(O^{\text{obs}}; \eta^*)\} = S_{a_0}(t_0).$$

We prove this condition case by case.

Case 1: $(S^*_Y, G^*_N) = (S_{0Y}, G_{0N})$**.** Under this condition, we have

$$\phi^F_{a_0,t_0}(\tilde{y}, \delta_y, x, n; \pi^*, G^*_Y) = S_{0Y}(t_0 \mid a_0, x, n) \left\{1 - \frac{I(a = a_0)}{\pi^*(a_0 \mid x)} H_{S_{0Y}, G^*_Y, t_0, a_0}(\tilde{y}, \delta_y, x, n)\right\},$$

where

$$\begin{aligned} H_{S_Y, G_Y, t_0, a_0}(\tilde{y}, \delta_y, x, n) = &\frac{I(\tilde{y} \le t_0, \delta_y = 1)}{S_Y(\tilde{y} \mid a_0, x, n) G_Y(\tilde{y} \mid a_0, x)} \\ &- \int_0^{\min(t_0, \tilde{y})} \frac{\Lambda_Y(du \mid a_0, x, n)}{S_Y(u \mid a_0, x, n) G_Y(u \mid a_0, x)}. \end{aligned} \tag{D.1}$$

We note that $\mathbb{E}\{H_{S_{0Y},G_Y^*,t_0,a_0}(\tilde{Y},\Delta_Y,X,N) \mid X=x,N=n\}$ equals

$$\int_0^{t_0} \frac{S_{0Y}(y- \mid a_0,x,n)G_{0Y}(y \mid a_0,x)}{S_{0Y}(y \mid a_0,x,n)G_Y^*(y \mid a_0,x)}\Lambda_{0Y}(dy \mid a_0,x,n)$$
$$-\int_0^{t_0} \frac{S_{0Y}(u- \mid a_0,x,n)G_{0Y}(u \mid a_0,x)}{S_{0Y}(u \mid a_0,x,n)G_Y^*(u \mid a_0,x)}\Lambda_{0Y}(du \mid a_0,x,n)=0,$$

then the full-data influence function is unbiased:

$$\mathbb{E}\{\phi^F_{a_0,t_0}(O^{\text{full}};\pi^*,G_Y^*)\}=\mathbb{E}\{S_{0Y}(t_0 \mid a_0,X,N)\}=S_{a_0}(t_0).$$

In addition, for the first term in $\phi^{\text{obs}}_{a_0,t_0}$, we have

$$\begin{aligned}
& \mathbb{E}\left\{\frac{\Delta_N}{G_{0N}(\tilde{N} \mid a_0,X,\tilde{Y},\Delta_Y)}\phi^F_{a_0,t_0} \mid \tilde{Y},\Delta_Y,a_0,X,N\right\}\\
= \;& \mathbb{E}\left\{\frac{I(N\le C_N)}{G_{0N}(\tilde{N} \mid a_0,X,\tilde{Y},\Delta_Y)}\phi^F_{a_0,t_0} \mid \tilde{Y},\Delta_Y,a_0,X,N\right\}\\
= \;& \phi^F_{a_0,t_0}(\tilde{Y},\Delta_Y,a_0,X,N)\cdot\frac{\mathbb{E}(C_N\ge N \mid \tilde{Y},\Delta_Y,a_0,X,N)}{G_{0N}(\tilde{N} \mid a_0,X,\tilde{Y},\Delta_Y)}\\
= \;& \phi^F_{a_0,t_0}(\tilde{Y},\Delta_Y,a_0,X,N).
\end{aligned}$$

Thus, $\mathbb{E}\{\Delta_N\phi^F_{a_0,t_0}/G_{0N}(\tilde{N} \mid a_0,X,\tilde{Y},\Delta_Y)\}=S_{a_0}(t_0)$.

As for the second term,

$$\begin{aligned}
& \mathbb{E}\left\{\frac{1-\Delta_N}{G_{0N}(\tilde{N} \mid a_0,X,\tilde{Y},\Delta_Y)}Q^*(\tilde{N})\right\}\\
= \;& \mathbb{E}\left\{\frac{I(C_N<N)}{G_{0N}(C_N \mid a_0,X,\tilde{Y},\Delta_Y)}Q^*(C_N)\right\}\\
= \;& \mathbb{E}\left\{\int_0^N \frac{Q^*(u)}{G_{0N}(u \mid a_0,X,\tilde{Y},\Delta_Y)}G_{0N}(u \mid a_0,X,\tilde{Y},\Delta_Y)d\Lambda_{C_N}(u \mid a_0,X,\tilde{Y},\Delta_Y)\right\}\\
= \;& \mathbb{E}\left\{\int_0^N Q^*(u)d\Lambda_{C_N}(u \mid a_0,X,\tilde{Y},\Delta_Y)\right\}.
\end{aligned}$$

The third term equals to

$$\begin{aligned}
& \mathbb{E}\left\{\int_0^\infty I(N\ge u)I(C_N\ge u)\frac{Q^*(u)}{G_{0N}(u \mid a_0,X,\tilde{Y},\Delta_Y)}d\Lambda_{C_N}(u \mid a_0,X,\tilde{Y},\Delta_Y)\right\}\\
= \;& \mathbb{E}\left[\int_0^\infty \mathbb{E}\{I(N\ge u)I(C_N\ge u) \mid \tilde{Y},\Delta_Y,a_0,X,N\}\frac{Q^*(u)}{G_{0N}(u \mid a_0,X,\tilde{Y},\Delta_Y)}d\Lambda_{C_N}(u \mid a_0,X,\tilde{Y},\Delta_Y)\right]\\
= \;& \mathbb{E}\left\{\int_0^N Q^*(u)d\Lambda_{C_N}(u \mid a_0,X,\tilde{Y},\Delta_Y)\right\}
\end{aligned}$$

Finally, under Case 1 (i.e.,$(S_Y^*, G_N^*) = (S_{0Y}, G_{0N})$), we have $E_0\{\phi^{\text{obs}}_{a_0,t_0}(O^{\text{obs}}; \pi^*, G_Y^*, f_N^*)\} = S_{a_0}(t_0)$.

Case 2: $(S_Y^*, f_N^*) = (S_{0Y}, f_{0N})$**.** Again, $S_Y^* = S_{0Y}$ implies that the full-data influence function is unbiased:

$$\mathbb{E}\{\phi^F_{a_0,t_0}(O^{\text{full}}; \pi^*, G_Y^*)\} = S_{a_0}(t_0).$$

The first term in $\phi^{\text{obs}}_{a_0,t_0}$ becomes

$$\begin{aligned} & \mathbb{E}\left[\mathbb{E}\left\{\frac{\Delta_N}{G_N^*(\tilde{N} \mid a_0, X, \tilde{Y}, \Delta_Y)}\phi^F_{a_0,t_0} \mid \tilde{Y}, \Delta_Y, a_0, X, N\right\}\right] \\ = & \mathbb{E}\left[\mathbb{E}\left\{\frac{I(N \le C_N)}{G_N^*(N \mid a_0, X, \tilde{Y}, \Delta_Y)}\phi^F_{a_0,t_0} \mid \tilde{Y}, \Delta_Y, a_0, X, N\right\}\right] \\ = & \mathbb{E}\left\{\frac{G_{0N}(N \mid a_0, X, \tilde{Y}, \Delta_Y)}{G_N^*(\tilde{N} \mid a_0, X, \tilde{Y}, \Delta_Y)}\phi^F_{a_0,t_0}\right\}. \end{aligned}$$

Since $f_N^* = f_{0N}$, we define

$$Q_0^*(u \mid \tilde{Y}, \Delta_Y, a_0, X) \equiv \mathbb{E}\left\{\phi^F_{a_0,t_0}(\tilde{Y}, \Delta_Y, X, N; \pi^*, G_Y^*, G_N^*, S_{0Y}) \mid N \ge u, \tilde{Y}, \Delta_Y, a_0, X\right\}.$$

The second term in $\phi^{\text{obs}}_{a_0,t_0}$ becomes

$$\mathbb{E}\left\{\int_0^N \frac{Q_0^*(u)}{G_N^*(u \mid a_0, X, \tilde{Y}, \Delta_Y)} G_{0N}(u \mid a_0, X, \tilde{Y}, \Delta_Y) d\Lambda_{C_N}(u \mid a, X, \tilde{Y}, \Delta_Y)\right\},$$

while the third term becomes

$$\mathbb{E}\left\{\int_0^N \frac{Q_0^*(u)}{G_N^*(u \mid a_0, X, \tilde{Y}, \Delta_Y)} G_{0N}(u \mid a_0, X, \tilde{Y}, \Delta_Y) d\Lambda^*_{C_N}(u \mid a, X, \tilde{Y}, \Delta_Y)\right\}.$$

Define $W(u \mid a, x, \tilde{y}, \delta_y) \equiv G_{0N}(u \mid a, x, \tilde{y}, \delta_y)/G_N^*(u \mid a, x, \tilde{y}, \delta_y)$, we have

$$dW(u \mid a_0, x, \tilde{y}, \delta_y) = W(u \mid a_0, x, \tilde{y}, \delta_y)\{d\Lambda^*_{C_N}(u \mid a_0, x, \tilde{y}, \delta_y) - d\Lambda_{C_N}(u \mid a_0, x, \tilde{y}, \delta_y)\}.$$

Thus,

$$\begin{aligned} & \mathbb{E}\left\{\int_0^N \frac{Q_0^*(u)}{G_N^*(u \mid a_0, X, \tilde{Y}, \Delta_Y)} G_{0N}(u \mid a_0, X, \tilde{Y}, \Delta_Y) d\Lambda_{C_N}(u \mid a, X, \tilde{Y}, \Delta_Y)\right\} \\ & -\mathbb{E}\left\{\int_0^N \frac{Q_0^*(u)}{G_N^*(u \mid a_0, X, \tilde{Y}, \Delta_Y)} G_{0N}(u \mid a_0, X, \tilde{Y}, \Delta_Y) d\Lambda^*_{C_N}(u \mid a, X, \tilde{Y}, \Delta_Y)\right\} \\ = & -\mathbb{E}\left\{\int_0^N Q_0^*(u) dW(u \mid a_0, X, \tilde{Y}, \Delta_Y)\right\} \end{aligned}$$

Further, we have

$$\begin{aligned}
& \mathbb{E}\left\{\int_0^N Q_0^*(u)dW(u \mid a_0, X, \tilde{Y}, \Delta_Y) \mid \tilde{Y}, \Delta_Y, a_0, X\right\} \\
= & \int_0^\infty \mathbb{E}\{I(N \geq u)Q_0^*(u) \mid \tilde{Y}, \Delta_Y, a_0, X\}dW(u \mid a_0, X, \tilde{Y}, \Delta_Y) \\
= & \int_0^\infty \mathbb{E}\{I(N \geq u)\phi_{a_0,t_0}^F \mid \tilde{Y}, \Delta_Y, a_0, X\}dW(u \mid a_0, X, \tilde{Y}, \Delta_Y).
\end{aligned}$$

In addition, from $W(N \mid a_0, X, \tilde{Y}, \Delta_Y) = 1 + \int_0^\infty I(N \geq u)dW(u \mid a_0, X, \tilde{Y}, \Delta_Y)$, we have

$$\begin{aligned}
& \mathbb{E}\{W(N \mid a_0, X, \tilde{Y}, \Delta_Y)\phi_{a_0,t_0}^F \mid \tilde{Y}, \Delta_Y, a_0, X\} \\
= & \mathbb{E}\{\phi_{a_0,t_0}^F \mid \tilde{Y}, \Delta_Y, a_0, X\} + \int_0^\infty \mathbb{E}\{I(N \geq u)\phi_{a_0,t_0}^F \mid \tilde{Y}, \Delta_Y, a_0, X\}dW(u \mid a_0, X, \tilde{Y}, \Delta_Y).
\end{aligned}$$

Combining the above displays, conditional on $(\tilde{Y}, \Delta_Y, A = a_0, X)$,

$$\begin{aligned}
& \mathbb{E}\left[\left\{W(N \mid a_0, X, \tilde{Y}, \Delta_Y)\phi_{a_0,t_0}^F - \int_0^N Q_0^*(u)dW(u \mid a_0, X, \tilde{Y}, \Delta_Y)\right\} \mid \tilde{Y}, \Delta_Y, a_0, X\right] \\
= & \mathbb{E}\{\phi_{a_0,t_0}^F \mid \tilde{Y}, \Delta_Y, a_0, X\}.
\end{aligned}$$

Taking expectation on both sides yields

$$\begin{aligned}
\mathbb{E}\{\phi_{a_0,t_0}^{\text{obs}}(O^{\text{obs}}; \pi^*, G_Y^*, G_N^*)\} &= \mathbb{E}\left[W(N \mid a_0, X, \tilde{Y}, \Delta_Y)\phi_{a_0,t_0}^F - \int_0^N Q_0^*(u)dW(u \mid a_0, X, \tilde{Y}, \Delta_Y)\right] \\
&= \mathbb{E}\{\phi^F(O^{\text{full}}; \pi^*, G_Y^*)\} = S_{a_0}(t_0).
\end{aligned}$$

Case 3: $(\pi^*, G_Y^*, G_N^*) = (\pi_0, G_{0Y}, G_{0N})$**..** Under this condition, we have

$$\phi_{a_0,t_0}^F(\tilde{y}, \delta_y, x, n; S^*) = S_Y^*(t_0 \mid a_0, x, n)\left\{1 - \frac{I(a = a_0)}{\pi_0(a_0 \mid x)}H_{S_Y^*, G_{0Y}, t_0, a_0}(\tilde{y}, \delta_y, x, n)\right\},$$

where $H_{S_Y, G_Y, t_0, a_0}(\tilde{y}, \delta_y, x, n)$ is defined in (D.1). We note that $\mathbb{E}\{H_{S_Y^*, G_{0Y}, t_0, a_0}(\tilde{Y}, \Delta_Y, X, N) \mid X = x, N = n\}$ equals to

$$\begin{aligned}
& \int_0^{t_0} \frac{S_{0Y}(y- \mid a_0, x, n)G_{0Y}(y \mid a_0, x)}{S_Y^*(y \mid a_0, x, n)G_{0Y}(y \mid a_0, x)}\Lambda_{0Y}(dy \mid a_0, x, n) \\
- & \int_0^{t_0} \frac{S_{0Y}(u- \mid a_0, x, n)G_{0Y}(u \mid a_0, x)}{S_Y^*(u \mid a_0, x, n)G_{0Y}(u \mid a_0, x)}\Lambda_Y^*(du \mid a_0, x, n) \\
= & \int_0^{t_0} \frac{S_{0Y}(u- \mid a_0, x, n)}{S_Y^*(u \mid a_0, x, n)}(\Lambda_{0Y} - \Lambda_Y^*)(du \mid a_0, x, n).
\end{aligned}$$

The Duhamel equation applied to $S_Y^* - S_{0Y}$ gives

$$S_Y^*(t \mid a_0, x, n) - S_{0Y}(t \mid a_0, x, n) = S_Y^*(t \mid a_0, x, n) \int_0^t \frac{S_{0Y}(u- \mid a_0, x, n)}{S_Y^*(u \mid a_0, x, n)} (\Lambda_{0Y} - \Lambda_Y^*)(du \mid a_0, x, n).$$

Thus, $\mathbb{E}\{H_{S_Y^*, G_{0Y}, t_0, a_0}(\tilde{Y}, \Delta_Y, X, N) \mid X = x, N = n\}$ equals to

$$\frac{S_Y^*(t_0 \mid a_0, x, n) - S_{0Y}(t_0 \mid a_0, x, n)}{S_Y^*(t_0 \mid a_0, x, n)}.$$

Therefore,

$$\begin{aligned} & \mathbb{E}\{\phi_{a_0,t_0}^F(\tilde{Y}, \Delta_Y, X, N; S_Y^*) \mid X = x, N = n\} \\ = \;& S_Y^*(t_0 \mid a_0, x, n) \left[1 - \frac{S_Y^*(t_0 \mid a_0, x, n) - S_{0Y}(t_0 \mid a_0, x, n)}{S_Y^*(t_0 \mid a_0, x, n)}\right] \\ = \;& S_{0Y}(t_0 \mid a_0, x, n). \end{aligned}$$

Taking expectation over (X, N) yields

$$\mathbb{E}\{\phi_{a_0,t_0}^F(O^{\text{full}}; S_Y^*)\} = S_{a_0}(t_0).$$

Since $G_N^* = G_{0N}$, we have proved in Case 1 that

$$\mathbb{E}\{\phi_{a_0,t_0}^{\text{obs}}(O^{\text{obs}}; f_N^*, S_Y^*)\} = \mathbb{E}\{\phi_{a_0,t_0}^F(O^{\text{full}}; S_Y^*)\} = S_{a_0}(t_0).$$

Case 4: $(\pi^*, G_Y^*, f_N^*) = (\pi_0, G_{0Y}, f_{0N})$**.** We have proved in Case 3 that

$$\mathbb{E}\{\phi_{a_0,t_0}^F(O^{\text{full}}; S_Y^*)\} = S_{a_0}(t_0).$$

In addition, when $f_N^* = f_{0N}$, from Case 2, we have shown that

$$\mathbb{E}\{\phi_{a_0,t_0}^{\text{obs}}(O^{\text{obs}}; S_Y^*, G_N^*)\} = E\{\phi_{a_0,t_0}^F(O^{\text{full}}; S_Y^*)\} = S_{a_0}(t_0).$$

Combining the four cases, the observed-data efficient estimator is consistent whenever one of

$$\begin{aligned} & (S_Y^*, G_N^*) = (S_{0Y}, G_{0N}), \qquad (S_Y^*, f_N^*) = (S_{0Y}, f_{0N}), \\ & (\pi^*, G_Y^*, G_N^*) = (\pi_0, G_{0Y}, G_{0N}), \qquad (\pi^*, G_Y^*, f_N^*) = (\pi_0, G_{0Y}, f_{0N}). \end{aligned}$$

holds. This proves the claimed multiple robustness.

D.4 Proof of Theorem 3

With the same decomposition as in the proof of Theorem 2,

$$\widehat{S}_{n,a_0}(t_0) - S_{a_0}(t_0) = R_{1n} + R_{2n} + R_{3n},$$

where

$$\begin{aligned} R_{1n} &= (\mathbb{P}_n - \mathbb{E})\phi^{\text{obs}}_{a_0,t_0}, \\ R_{2n} &= \frac{1}{K}\sum_{k=1}^{K} \frac{Kn_k}{n}(\mathbb{P}_{n,k} - \mathbb{E})(\widehat{\phi}^{\text{obs}}_{n,k,a_0,t_0} - \phi^{\text{obs}}_{a_0,t_0}), \\ R_{3n} &= \frac{1}{K}\sum_{k=1}^{K} \frac{Kn_k}{n}\mathbb{E}(\widehat{\phi}^{\text{obs}}_{n,k,a_0,t_0} - \phi^{\text{obs}}_{a_0,t_0}). \end{aligned}$$

We have shown in the proof of Theorem 2 that

$$R_{2n} = o_p(n^{-1/2}).$$

For the third term R_{3n}, we analyze the three terms in the observed-data EIF separately.

For the first term, define $\widehat{W}_{n,k,N}(u \mid a, x, \widetilde{y}, \delta_y) \equiv G_{0N}(u \mid a, x, \widetilde{y}, \delta_y)/\widehat{G}_{n,k,N}(u \mid a, x, \widetilde{y}, \delta_y)$, then

$$\mathbb{E}\left\{ \frac{\Delta_N}{\widehat{G}_{n,k,N}(\widetilde{N} \mid a_0, X, \widetilde{Y}, \Delta_Y)} \widehat{\phi}^F_{n,k,a_0,t_0} \right\} = \mathbb{E}\{\widehat{W}_{n,k,N}(N \mid a_0, X, \widetilde{Y}, \Delta_Y)\widehat{\phi}^F_{n,k,a_0,t_0}\}.$$

Using $\widehat{W}_{n,k,N}(N \mid a_0, X, \widetilde{Y}, \Delta_Y) = 1 + \int_0^\infty I(N \geq u) d\widehat{W}_{n,k,N}(u \mid a_0, X, \widetilde{Y}, \Delta_Y)$, we obtain

$$\begin{aligned} & \mathbb{E}\left\{ \frac{\Delta_N}{\widehat{G}_{n,k,N}(\widetilde{N} \mid a_0, X, \widetilde{Y}, \Delta_Y)} \widehat{\phi}^F_{n,k,a_0,t_0} \right\} \\ = \; & \mathbb{E}\left\{ \int_0^N \widehat{\phi}^F_{n,k,a_0,t_0} d\widehat{W}_{n,k,N}(u \mid a_0, X, \widetilde{Y}, \Delta_Y) \right\} + \mathbb{E}(\widehat{\phi}^F_{n,k,a_0,t_0}) \\ = \; & \mathbb{E}\left\{ \int_0^N \widehat{\phi}^F_{n,k,a_0,t_0} d\widehat{W}_{n,k,N}(u \mid a_0, X, \widetilde{Y}, \Delta_Y) \right\} + S_{a_0}(t_0) + o_p(n^{-1/2}), \end{aligned}$$

where the second equality above is by Theorem 3 of[3].

Next consider the second and third terms:

$$\begin{aligned} & \mathbb{E}\left\{ \frac{1 - \Delta_N}{\widehat{G}_{n,k,N}(\widetilde{N} \mid a_0, X, \widetilde{Y}, \Delta_Y)} \widehat{Q}_{n,k}(\widetilde{N} \mid a_0, X, \widetilde{Y}, \Delta_Y) \right\} \\ - \; & \mathbb{E}\left\{ \int_0^{\widetilde{N}} \frac{\widehat{Q}_{n,k}(u \mid a_0, X, \widetilde{Y}, \Delta_Y)}{\widehat{G}_{n,k,N}(u \mid a_0, X, \widetilde{Y}, \Delta_Y)} d\widehat{\Lambda}_{n,k,C_N}(u \mid a_0, X, \widetilde{Y}, \Delta_Y) \right\}. \end{aligned}$$

By

$$d\widehat{W}_{n,k,N}(u \mid a_0, x, \widetilde{y}, \delta_y) = \widehat{W}_{n,k,N}(u \mid a_0, x, \widetilde{y}, \delta_y)\{d\widehat{\Lambda}_{n,k,C_N}(u \mid a_0, x, \widetilde{y}, \delta_y) - d\Lambda_{0C_N}(u \mid a_0, x, \widetilde{y}, \delta_y)\},$$

the above difference equals

$$-\mathbb{E}\left\{\int_0^N \widehat{Q}_{n,k}(u \mid a_0, X, \widetilde{Y}, \Delta_Y) d\widehat{W}_{n,k,N}(u \mid a_0, X, \widetilde{Y}, \Delta_Y)\right\}.$$

Combining with the expansion of the first term:

$$\begin{aligned} & \mathbb{E}\left\{\int_0^N \widehat{\phi}^F_{n,k,a_0,t_0} d\widehat{W}_{n,k,N}(u \mid a_0, X, \widetilde{Y}, \Delta_Y)\right\} \qquad \text{(D.2)} \\ - \quad & \mathbb{E}\left\{\int_0^N \widehat{Q}_{n,k}(u \mid a_0, X, \widetilde{Y}, \Delta_Y) d\widehat{W}_{n,k,N}(u \mid a_0, X, \widetilde{Y}, \Delta_Y)\right\} \\ = \quad & \mathbb{E}\left[\int_0^N \{\widehat{\phi}^F_{n,k,a_0,t_0} - Q_k(u)\} d\widehat{W}_{n,k,N}(u)\right] + \mathbb{E}\left[\int_0^N \{Q_k(u) - \widehat{Q}_{n,k}(u)\} d\widehat{W}_{n,k,N}(u)\right], \qquad \text{(D.3)} \end{aligned}$$

where $Q_k(u \mid a, x, \widetilde{y}, \delta_y) = \mathbb{E}(\widehat{\phi}^F_{n,k,a_0,t_0} \mid N \geq u, a, x, \widetilde{y}, \delta_y)$. Since

$$\begin{aligned} & \mathbb{E}\left[\int_0^N \{\widehat{\phi}^F_{n,k,a_0,t_0} - Q_k(u)\} d\widehat{W}_{n,k,N}(u) \mid a_0, X, \widetilde{Y}, \Delta_Y\right] \\ = \quad & \int_0^\infty \mathbb{E}\left[I(N \geq u)\{\widehat{\phi}^F_{n,k,a_0,t_0} - Q_k(u)\} \mid a_0, X, \widetilde{Y}, \Delta_Y\right] d\widehat{W}_{n,k,N}(u) \\ = \quad & \int_0^\infty [P(N \geq u \mid a_0, X, \widetilde{Y}, \Delta_Y) Q_k(u) - \mathbb{E}\{I(N \geq u) Q_k(u) \mid a_0, X, \widetilde{Y}, \Delta_Y\}] d\widehat{W}_{n,k,N}(u) \\ = \quad & 0, \end{aligned}$$

equation (D.2) reduces to

$$\begin{aligned} & \mathbb{E}\left[\int_0^N \{Q_k(u) - \widehat{Q}_{n,k}(u)\} d\widehat{W}_{n,k,N}(u)\right] \\ = \quad & \mathbb{E}\left[\int_0^N \{Q_k(u) - \widehat{Q}_{n,k}(u)\} \widehat{W}_{n,k,N}(u)\{d\widehat{\Lambda}_{n,k,C_N}(u) - d\Lambda_{0C_N}.(u)\}\right] \end{aligned}$$

By Condition C2, we have $\widehat{W}_{n,k,N}$ is bounded. In addition, by condition C4,

$$\mathbb{E}\left[\int_0^N \{Q_k(u) - \widehat{Q}_{n,k}(u)\} d\widehat{W}_{n,k,N}(u)\right] = o_p(n^{-1/2}).$$

Finally, we obtain

$$R_{3n} = S_{a_0}(t_0) + o_p(n^{-1/2}) - \mathbb{E}\phi^{\text{obs}}_{a_0,t_0} = o_p(n^{-1/2}).$$

Therefore,

$$\widehat{S}_{n,a_0}(t_0) - S_{a_0}(t_0) = (\mathbb{P}_n - \mathbb{E})\phi^{\text{obs}}_{a_0,t_0} + o_p(n^{-1/2}).$$

Multiplying by $\sqrt{n}$,

$$\sqrt{n}\left\{\widehat{S}_{n,a_0}(t_0) - S_{a_0}(t_0)\right\} = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\left\{\phi^{\text{obs}}_{a_0,t_0}(O_i) - S_{a_0}(t_0)\right\} + o_p(1).$$

The result follows from the central limit theorem.

D.5 Proposition 5

Recall our goal is to derive the relationship between the remaining nuisance parameter in the EIF $\phi^{\text{obs}}_{a_0,t_0}$ – $f_{0N}(n \mid a, x, \tilde{y}, \delta_y)$ – and the nuisances parameters compatible with a "one-way" factorization of the joint likelihood $(Y, N) \mid (A, X)$ – $f_N(n \mid a, x)$ and $S_{0Y}(y \mid a, x, n)$. This helps ensure that our nuisances are compatible with a joint model for $(Y, N) \mid (A, X)$. The key element of our derivation is an application of Bayes rule.

Specifically, we decompose the target nuisance parameter according to the event indicator δ_y.

$$\begin{aligned} f_{0N}(n \mid a, x, \tilde{y}, \delta_y) &= \delta_y f(n \mid a, x, c_y > \tilde{y}, y = \tilde{y}) + (1 - \delta_y) f_N(n \mid a, x, c_Y = \tilde{y}, y > \tilde{y}) \\ &= \delta_y f(n \mid a, x, y = \tilde{y}) + (1 - \delta_y) f(n \mid a, x, y > \tilde{y}) \end{aligned}$$

The cancellation of the terms $c_y > \tilde{y}$ and $c_y = \tilde{y}$ in the second line is the result of independent censoring (Assumption 4). Next, we apply Bayes rule to the constituent pieces.

$$\begin{aligned} f_{0N}(n \mid a, x, \tilde{y}, \delta_y) &= \delta_y \left(\frac{f_{0Y}(Y = \tilde{y} \mid a, x, n) f(n \mid a, x)}{f_Y(Y = \tilde{y} \mid a, x)} \right) + (1 - \delta_y) \frac{f(n \mid a, x) S_{0Y}(\tilde{y} \mid a, x, n)}{S_Y(\tilde{y} \mid a, x)} \\ &= f_N(n \mid a, x) \cdot \frac{L(\tilde{y}, \delta_y \mid a, x, n)}{L(\tilde{y}, \delta_y \mid a, x)} \end{aligned}$$

where

$$\frac{L(\tilde{y}, \delta_y \mid a, x, n)}{L(\tilde{y}, \delta_y \mid a, x)} = \left[\delta_y \left(\frac{f_{0Y}(\tilde{y} \mid a, x, n)}{f_Y(\tilde{y} \mid a, x)} \right) + (1 - \delta_y) \frac{S_{0Y}(\tilde{y} \mid a, x, n)}{S_Y(\tilde{y} \mid a, x)} \right].$$

$S_Y(u \mid a, x) = \int_0^\infty S_{0Y}(u \mid a, x, n) f_N(n \mid a, x) dn$ by Tower law. f_{0Y} and f_Y refer to the conditional density functions associated with S_{0Y} and S_Y, which can be obtained from the survival functions using the following identity $f(x) = -(d/du)S(u) \mid_{u=x}$.

Hence, a compatible estimator of $f_{0N}(n \mid a, x, \tilde{y}, \delta_y)$ can be obtained from the nuisances that define a one-way factorization of the joint likelihood: $f_N(n \mid a, x)$ and $S_{0Y}(y \mid a, x, n)$. □

Bibliography

[1] Paul R. Rosenbaum. The consequences of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society: Series B*, 147, 1984.

[2] CE Frangakis and DB Rubin. Principal stratification in causal inference. *Biometrics*, 58:21–29, 2002.

[3] Ted Westling, Alex Luedtke, Peter B. Gilbert, and Marco Carone. Inference for treatment-specific survival curves using machine learning. *Journal of the American Statistical Association*, 2023.

[4] Peter J. Bickel, Chris A.J. Klaassen, Ya'acov Ritov, and Jon A. Wellner. *Efficient and adaptive estimation for semiparametric models*. Johns Hopkins University Press Baltimore, 1993.

[5] Anastasios A Tsiatis. *Semiparametric Theory and Missing Data*. Springer Series in Statistics. Springer, New York, NY, 2006 edition, June 2006.