

©Copyright 2025

Antonio Olivas-Martinez

Advances in Proximal Inference for Continuous Exposures,
Estimation of Ill-Posed Regression, and Non-Inferiority Assessment
in Active-Controlled Trials

Antonio Olivas-Martinez

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

Doctor of Philosophy

University of Washington

2025

Reading Committee:

Andrea Rotnitzky, Chair

Holly Janes

Peter Gilbert

Program Authorized to Offer Degree:

Biostatistics

University of Washington

Abstract

Advances in Proximal Inference for Continuous Exposures, Estimation of Ill-Posed Regression, and Non-Inferiority Assessment in Active-Controlled Trials

Antonio Olivas-Martinez

Chair of the Supervisory Committee:

Andrea Rotnitzky

Department of Biostatistics

This dissertation advances proximal inference methods for continuous exposures and ill-posed regression problems, alongside non-inferiority assessment in active-controlled trials.

An introduction to the topics is provided in Chapter 1. Chapter 2 develops a unifying framework for non-inferiority evaluation in active-controlled trials, where placebo arms are unavailable, enabling systematic comparison of existing methods in terms of type I error, power, and robustness to transportability misspecifications. Chapter 3, motivated by the analysis of immune correlates of protection in COVID-19 vaccine trials, extends the proximal inference framework to identify and estimate the mean outcome under modified treatment policies, enabling causal analysis of continuous exposures in the presence of unmeasured confounding. Chapter 4 addresses the challenge of estimating ill-posed nuisance functions in proximal inference and related settings, developing a novel finite-sample analysis of kernel-based regularized adversarial stabilized estimators and establishing conditions under which debiased, influence-function-based one-step estimators for a broad class of estimands, achieve root-n consistency and asymptotic normality. Together, these contributions support more robust inference in complex causal problems with unmeasured confounding or design constraints.

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	vi
Chapter 1: Introduction	1
Chapter 2: A General Framework for Assessing Non-Inferiority in Active-Controlled Trials	4
2.1 Introduction	4
2.2 Motivational Example	6
2.3 Non-Inferiority Hypothesis and Identification Challenges	8
2.4 The General Framework	10
2.5 Operating Characteristics within the General Framework	16
2.6 Application to Designing an HIV Prevention Trial	19
2.7 Discussion	29
Chapter 3: Proximal Causal Inference for Modified Treatment Policies	31
3.1 Introduction	31
3.2 Motivating Example	34
3.3 Target of Inference and Identification	35
3.4 Estimation	45
3.5 Numerical Experiments	51
3.6 Application to the ENSEMBLE Trial	57
3.7 Discussion	60
Chapter 4: Source Condition-Based Analysis and Assumption Comparisons for RKHS-Norm Penalized Kernel Adversarial Stabilized Estimators	61

4.1	Introduction	61
4.2	Target Parameter	63
4.3	Debiased Estimation of $\theta(P_0)$	68
4.4	Adversarial Estimation of the Nuisance Functions	71
4.5	Finite-sample Bounds on the Errors of the KRAS Estimators	82
4.6	Contrasting Kernel based Adversarial Estimation Strategies	94
4.7	Discussion	105
Appendix A: Technical Details of Chapter 2		106
A.1	Formulas for Unconditional and Conditional Power	106
A.2	Detectable Design Alternatives and Maximum Unconditional Power	108
A.3	Acceptable Assumed Level of Non-Constancy and Tolerable Level of Non-Constancy	110
Appendix B: Supplementary Material of Chapter 3		112
B.1	Proximal Identification and Estimation in a potentially restricted population	112
B.2	Estimation of the Bridge Functions	120
B.3	Additional Numerical Experiments and Implementation	131
B.4	Proofs of Main Results	149
B.5	Supporting Lemmas	155
B.6	Existence of Bridge Functions for our Data Generating Mechanism	160
Appendix C: Supplementary Material of Chapter 4		168
C.1	Proofs of Mathematical Statements	168
C.2	Supporting Lemmas	176
C.3	Summary of Existing Regularized Adversarial Estimation Proposals	197

LIST OF FIGURES

Figure Number		Page
2.1	Maximum unconditional power for detecting design alternatives across our selected methods and criteria. SM, synthesis method; BA-SM, bias-adjusted synthesis method; OD, Odem-Davis method. The horizontal dashed line represents an unconditional power of 0.90. The vertical dotted line represents an experimental intervention with a PE of 95%.	20
3.1	A causal directed acyclic graph of proximal causal inference. A and Y represent the exposure and outcome of interest; L and U the observed and unobserved confounders; and Z and W the negative control treatments and negative control outcomes, respectively.	34
3.2	Boxplots of the estimators from numerical experiments with data generated under varying conditional correlations between the proxies Z and W and the unmeasured confounder U , with sample sizes $n = 750, 1500, 3000$, and 6000 . The dashed line represents the true parameter value. Non-proximal estimators rely on the (incorrect) assumption of no unmeasured confounding.	54
3.3	Ratio of empirical variance to the mean estimated asymptotic variance for each estimator, based on numerical experiments with data generated under varying conditional correlations between the proxies Z and W and the unmeasured confounder U , across sample sizes $n = 750, 1500, 3000$, and 6000 . The dashed line indicates a variance ratio of 1.	55
3.4	Coverage probabilities of the confidence intervals from numerical experiments with data generated under varying conditional correlations between the proxies Z and W and the unmeasured confounder U , across sample sizes $n = 750, 1500, 3000$, and 6000 . The dashed line indicates the nominal coverage level of 0.95. Non-proximal estimators rely on the (incorrect) assumption of no unmeasured confounding.	56

3.5	Results from application to data from the ENSEMBLE trial. (Panel A) Observed and counterfactual probabilities of symptomatic COVID-19 infection and (Panel B) observed and counterfactual vaccine efficacy (VE) against symptomatic COVID-19 infection for policies that shift everyone's observed antibody levels by 0 (no shift), 0.2, 0.4, 0.6, 0.8, and 1 logarithm. The bars indicate the 95% confidence interval at each shift without any correction.	59
B.1	Boxplots of the parametric estimators from numerical experiments with data generated under varying conditional correlations of the proxies Z and W with the unmeasured confounder U and sample sizes $n = 750, 3000, 12000$, and 48000. Dashed lines indicates the true parameter value.	138
B.2	Empirical coverage of 95% confidence intervals from numerical experiments with data generated under varying conditional correlations of the proxies Z and W with the unmeasured confounder U and sample sizes $n = 750, 3000, 12000$, and 48000. The dashed line indicates the nominal coverage level of 0.9.	139
B.3	Finite-sample performance of $\hat{\psi}_{CF}^{DR}$ and competitors based on numerical experiments <i>without</i> unmeasured confounding. Panel A: boxplots of the estimator across simulations. Panel B: ratio of the empirical variance to the mean estimated asymptotic variance. Panel C: empirical coverage of 95% confidence intervals. The dashed line indicates the true parameter value (Panel A), a variance ratio of 1 (Panel B), and nominal 95% coverage (Panel C).	141
B.4	Finite-sample performance of $\hat{\psi}_{CF}^{DR}$ for estimating the counterfactual MTP mean in a <i>restricted population</i> , based on numerical experiments with data generated under unmeasured confounding. Panel A: boxplots of the estimator across simulations. Panel B: ratio of the empirical variance to the mean estimated asymptotic variance. Panel C: empirical coverage of 95% confidence intervals. The dashed line indicates the true parameter value (Panel A), a variance ratio of 1 (Panel B), and nominal 95% coverage (Panel C).	142
B.5	Finite-sample performance of $\hat{\psi}_{CF}^{DR}$ across varying levels of conditional correlation between the proxies Z and W and U . Panel A: boxplots of the estimator across simulations. Panel B: ratio of the empirical variance to the mean estimated asymptotic variance. Panel C: empirical coverage of 95% confidence intervals. The dashed line indicates the true parameter value (Panel A), a variance ratio of 1 (Panel B), and nominal 95% coverage (Panel C).	144

B.6	Boxplots of the <i>correctly specified</i> parametric estimators from numerical experiments under a more challenging data generated-generating mechanism, where Z directly affects A and W directly affects Y . Results are shown across varying conditional correlations of the proxies Z and W with the unmeasured confounder U , and sample sizes $n = 750, 3000, 12000$, and 48000 . Dashed lines indicate the true parameter values.	166
B.7	Empirical coverage of 95% confidence intervals of <i>correctly specified</i> parametric estimators from numerical experiments under a more challenging data generated-generating mechanism, where Z directly affects A and W directly affects Y . Results are shown across varying conditional correlations of the proxies Z and W with the unmeasured confounder U , and sample sizes $n = 750, 3000, 12000$, and 48000 . Dashed lines indicate the true parameter values. The dashed line indicates the nominal coverage level of 0.9.	167

LIST OF TABLES

Table Number	Page
2.1 Common methods for assessing non-inferiority and their robustness to type I error control under constancy	13
2.2 Trial design specifications and operating characteristics for detecting 95% prevention efficacy with 90% power under the 50% effect preservation criterion, across methods and design strategies.	27
2.3 Trial design specifications and operating characteristics for detecting 95% prevention efficacy with 90% power under the 30% prevention efficacy criterion, across methods and design strategies.	28
C.1 Existing Adversarial Estimation Approaches	198
C.2 Comparison of Assumptions for Finite-Sample Bounds	199

ACKNOWLEDGMENTS

This dissertation is the result of many conversations, moments of doubt, generous acts of support, and sustained mentorship. I am profoundly thankful to those who stood by me throughout—challenging my ideas, encouraging my growth, and reminding me why this endeavor matters.

Above all, I owe my deepest gratitude to my advisor, Dr. Andrea Rotnitzky. Her unwavering commitment to my development as a researcher has shaped every stage of this dissertation. I am especially grateful for the time and care she devoted to teaching me how to write more clearly and effectively in statistics, and for encouraging me to engage more deeply with the theoretical foundations of my work. Her passion for research with real-world impact, and her guidance on how to ask meaningful questions and pursue them through coherent and well-structured lines of inquiry, have been instrumental in my growth as a scientist. Her mentorship has left a lasting imprint on how I think about and carry out research.

I am also sincerely thankful to the members of my dissertation committee—Dr. Holly Janes, Dr. Peter Gilbert, Dr. Fei Gao, and Dr. Carey Farquhar—for their thoughtful feedback, constructive critiques, and generous investment of time. Their insights and perspectives have meaningfully enriched this work. I am especially thankful to Dr. Holly Janes for her guidance and encouragement in the early phase of my dissertation work. Her example of combining methodological innovation with clarity

of communication continues to influence my approach to research, and I look forward to carrying those lessons forward.

I am grateful to the faculty and staff at the University of Washington for fostering an intellectually vibrant and collaborative academic environment. I am especially thankful to Dr. Lurdes Inoue, who taught me how to conduct applied collaborative research with clinicians and has consistently offered wise, practical advice for my career. Her guidance and genuine interest in my growth have meant a great deal. I am also thankful to Dr. Ruth Etzioni, whose encouragement and openness to collaboration early in my doctoral training broadened my perspective on applied statistical research and left a lasting impression on me. I would also like to thank Minh Vo, who was incredibly helpful in navigating the logistics of graduate school and who, like Lurdes, always made me feel seen and supported.

To my family—thank you. To my mother, Sofia: your unconditional love and unwavering belief in me have carried me through every high and low of this journey. To my late father, Juan: I carry your memory with pride and deep gratitude. To my brothers, Miguel and Oscar, and to the rest of my family—thank you for your love, encouragement, and patience throughout this long process.

I also feel incredibly lucky to have been surrounded by a wonderful community of peers and friends. To my cohort—Jess, Leah, Nina, Daniel, and Seth—thank you for your camaraderie, honesty, and steady support. To other friends in my program and beyond—Ameer, James, Jordan, Pawel, Zack, Arianna, Juandalyn, Elizabeth, Hangjun—thank you for the good conversations, shared frustrations, and steady encouragement along the way.

To my friends beyond academia—too many to list—thank you for grounding me,

reminding me there is life beyond grad school, and for your love, laughter, and enduring presence throughout it all.

To all of you—thank you

DEDICATION

To my mom, for her endless support, and to my father, in loving memory.

Chapter 1

INTRODUCTION

Active-controlled trials with non-inferiority objectives are well-suited for phase 2b/3 studies when existing interventions are effective and available, but new options offer advantages (Rothmann et al., 2011; Fleming et al., 2011; FDA et al., 2016). Traditional non-inferiority criteria require new interventions to preserve a high proportion of the active control’s effect. However, when the control is highly effective, alternative criteria based on sufficient efficacy versus a hypothetical placebo may be appropriate. A key challenge in these trials is the absence of a placebo arm, requiring reliance on historical data to estimate the active control effect and making assumptions about the extent to which this effect transports to the target population. A wide variety of methods exist for analyzing data from a non-inferiority trial (Wang et al., 2002; Hung et al., 2003, 2007; Snapinn and Jiang, 2008; Brittain et al., 2012; Odem-Davis and Fleming, 2013), and different non-inferiority criteria may be considered (Hung et al., 2003; Fleming, 2008; FDA et al., 2016; Fleming et al., 2021). An evaluation of the operating characteristics of the different analysis methods and non-inferiority criteria is lacking. In Chapter 2, we introduce a framework for non-inferiority assessment that encompasses all existing methods for analysis and accommodates traditional and alternative success criteria. Our framework facilitates the comparison of methods based on type I error, power, and robustness to misspecification of the transportability assumption. We apply this framework to design an HIV prevention trial with a highly effective active control.

Proximal causal inference has advanced the identification and inference of causal effects in scenarios involving unobserved confounding and two disjoint sets of observed strong proxies

for the unmeasured confounders—commonly referred to as negative control treatments and negative control outcomes (Miao et al., 2018; Tchetgen et al., 2020; Cui et al., 2024). It has predominantly been explored in the context of target estimands that contrast counterfactual worlds, where subjects undergo specific pre-specified treatment regimes, which may be randomized and could depend on observed covariates (Ghassami et al., 2022; Kallus et al., 2021). However, when dealing with continuous exposures, a compelling alternative estimand exists—the mean under a modified treatment regime (Haneuse and Rotnitzky, 2013; Díaz et al., 2023). This concept envisions a world in which each individual receives a treatment dose corresponding to their actual dose but modified in a pre-specified manner. In Chapter 3, we extend the proximal causal inference framework to identify and estimate the mean of an outcome under a modified treatment regime for continuous point-exposures, bridging an important gap in the literature. We illustrate our approach by applying it to immunobridging trials of COVID-19 vaccines, with focus on identifying correlates of protection.

In many applications, the target parameter depends on a nuisance function defined by a conditional moment restriction, whose estimation often leads to an ill-posed inverse problem (Santos, 2011; Severini and Tripathi, 2012; Escanciano and Li, 2013; Freyberger and Horowitz, 2015; Miao and Tchetgen Tchetgen, 2016; Miao et al., 2018; Tchetgen et al., 2020; Kallus et al., 2021; Ghassami et al., 2022; Bennett et al., 2022, 2023b; Li et al., 2023; Miao et al., 2024; Cui et al., 2024). Classical approaches, such as sieve-based GMM, approximate the restriction using a fixed set of test functions and may fail to capture important aspects of the solution (Newey and Powell, 2003; Ai and Chen, 2003, 2007). Adversarial estimators address this limitation by framing estimation as a game between an estimator and an adaptive critic (Bennett et al., 2019; Muandet et al., 2020b,a; Dikkala et al., 2020; Liao et al., 2020; Kallus et al., 2021; Ghassami et al., 2022; Mastouri et al., 2021; Bennett et al., 2022; Zhang et al., 2023; Bennett et al., 2023a,b; Park et al., 2024). In Chapter 4, we study the class of

Regularized Adversarial Stabilized (RAS) estimators that employ reproducing kernel Hilbert spaces (RKHSs) for both estimation and testing, with regularization via the RKHS norm. Our first contribution is a novel analysis that establishes finite-sample bounds for both the weak error and the root mean squared error of these estimators under interpretable source conditions, in contrast to existing results. Our second contribution is a detailed comparison of the assumptions underlying this RKHS-norm-regularized approach with those required for (i) RAS estimators using $\mathcal{L}^2$ penalties, and (ii) recently proposed, computationally stable Kernel Maximal Moment estimators.

Chapter 2

A GENERAL FRAMEWORK FOR ASSESSING NON-INFERIORITY IN ACTIVE-CONTROLLED TRIALS

2.1 *Introduction*

Active-controlled trials are the most frequently used phase 2b/3 trial design when a current intervention is known to be effective and is available, but new interventions are still warranted (Rothmann et al., 2011; Fleming et al., 2011; FDA et al., 2016). The new intervention may have superior or clinically-relevant efficacy and advantages in terms of reduced cost or toxicity, improved tolerability, adherence, or ease of implementation; or additional interventions may be needed to serve public health needs. In such contexts, a placebo-controlled design is likely not justified. In active-controlled trials, participants are randomized to one or more experimental interventions or to an active control, without a placebo arm. They may be designed to assess whether the new intervention is superior or non-inferior to the active control, where being non-inferior commonly is defined as preserving a fraction of the active control effect, often 50% (Fleming, 2008; FDA et al., 2016). Active-control trials designed to assess non-inferiority are often called non-inferiority trials.

A significant challenge in drawing inferences from active-controlled trials is the absence of a placebo arm. Assessing non-inferiority traditionally involves estimating the active control effect relative to placebo based on historical data, preferably from randomized trials, necessitating assumptions about the extent to which the historical estimate of active control effect applies to the target population. These assumptions must take into account possible effect modification by both measured and unmeasured modifiers, as well as factors such as

advancements in concomitant care, shifts in disease etiology or diagnostic criteria, evolving trial endpoints, and changes in the dose or regimen of the active control as its clinical use evolves (Fleming et al., 2011). The predominant assumption used is constancy, which essentially posits that the active control effect estimated in the historical trial(s) remains unchanged in the target population. Regrettably, any deviations from this assumption may undermine the validity of conclusions drawn from non-inferiority trials (Wang et al., 2002; Fleming et al., 2011; Rothmann et al., 2011; Odem-Davis and Fleming, 2013; FDA et al., 2016).

The primary methods for assessing non-inferiority based on active-controlled trials are the fixed-margin and synthesis methods (Wang et al., 2002; Rothmann et al., 2011; FDA et al., 2016). The fixed-margin method utilizes a predetermined margin to evaluate the efficacy of the experimental intervention against the active control, while the synthesis method integrates data from both the active-controlled trial and historical trial(s) in order to assess non-inferiority. Both methods have been extensively evaluated under constancy and non-constancy conditions. While the traditional synthesis method only controls type-I error under constancy, the fixed-margin method is demonstrated to be robust to deviations from constancy (Wang et al., 2002; Hung et al., 2003, 2007; Snapinn and Jiang, 2008; Brittain et al., 2012; Odem-Davis and Fleming, 2013). A generalization of the synthesis method (Odem-Davis and Fleming, 2013) improves robustness to non-constancy. A unified approach for comparing the operating characteristics of the various approaches is lacking.

Another challenge with non-inferiority design arises in the context of a highly effective active control. In such a context, an experimental intervention may still have public health impact even if it would not satisfy commonly employed preservation of effect non-inferiority criteria. For example, if the active control is 95% effective, a common 50% preservation of effect criterion would stipulate that a new intervention be more than 77.6% effective in

order to achieve success. This high bar may not be appropriate if the new intervention has advantages in terms of individual preference, cost, or feasibility of implementation. A compelling success criterion may be that the new intervention meets the success criterion that would have been used had a placebo-controlled design been possible. This is the spirit of a criterion proposed for supporting the design of COVID-19 vaccine non-inferiority trials (Fleming et al., 2021).

To address the aforementioned challenges, we propose a general framework for evaluating the non-inferiority of experimental interventions compared to active controls. Our framework accommodates traditional and new non-inferiority criteria and all existing methods for non-inferiority analysis, and facilitates systematic evaluation and comparison across methods based on operating characteristics that are either defined conditional on the historical data or unconditional. It also enables quantification of the degree of non-constancy accommodated by each method, and offers a design approach that formally factors in uncertainty in the active control effect based on historical data. We demonstrate the applicability of our framework by designing a future HIV prevention trial with a highly effective active control.

2.2 Motivational Example

Much success has been seen in recent years in biomedical prevention of HIV, with oral (Grant et al., 2010; Baeten et al., 2012; Thigpen et al., 2012; Van Damme et al., 2012; Marrazzo et al., 2015) and now injectable (Landovitz et al., 2021; Delany-Moretlwe et al., 2022; Bekker et al., 2024; Gilead Sciences, 2024) antiretrovirals proven effective for prevention, known as pre-exposure prophylaxis (PrEP). Long-acting Cabotegravir (CAB-LA) has proven highly effective and superior to oral PrEP in both women and men who have sex with men (Landovitz et al., 2021; Delany-Moretlwe et al., 2022). However, rollout, uptake, and adherence remain significant challenges, and additional interventions are clearly needed to achieve the

Joint United Nations Programme on HIV/AIDS (UNAIDS) targets for HIV incidence (UNAIDS, 2023; van der Straten et al., 2014; Henderson et al., 2023), including an effective HIV vaccine. HIV monoclonal antibodies, on-demand products, and alternative antiretrovirals and delivery devices are under investigation (Corey et al., 2021; Walker, 2021). In a setting where CAB-LA is licensed and available for use, a potential trial design to evaluate a novel antiretroviral as PrEP could be an active-controlled trial with CAB-LA as the active control, which seeks to establish that the novel antiretroviral is non-inferior to CAB-LA.

A key target population for such a trial would be women in sub-Saharan Africa, who remain at high risk of HIV despite advances in HIV prevention (UNAIDS, 2023; Murewanhema et al., 2022; Moyo et al., 2022). In this population, CAB-LA prevention efficacy (PE) have been recently estimated at 92.8% with a 95% confidence interval (CI): 76.1%–97.8% (Donnell et al., 2023). Given this high efficacy, the conventional 50% preservation of effect criterion may be unduly stringent. For instance, under this criterion, a new intervention would need to exceed a 73.2% PE to be deemed non-inferior—a threshold that could preclude promising interventions with important practical advantages. An alternative and arguably more appropriate criterion would be to assess whether the new intervention meets the same success threshold that would have applied in a placebo-controlled trial. In previous phase 3 placebo-controlled PrEP trials, this threshold corresponded to a minimum efficacy of 30% PE (Grant et al., 2010; Van Damme et al., 2012; Baeten et al., 2012; Marrazzo et al., 2015). A new antiretroviral that achieves at least 30% PE and offers advantages such as lower cost, improved adherence, reduced toxicity, or greater ease of implementation could have substantial public health impact. This example motivates the need for a flexible framework that accommodates such alternative non-inferiority criteria while maintaining rigorous statistical evaluation.

2.3 *Non-Inferiority Hypothesis and Identification Challenges*

We begin by introducing notation and framing the null hypothesis within our statistical framework. Let A be an intervention of interest, and let ψ_A denote a function characterizing the distribution of an outcome of interest in a population receiving intervention A . We define the effect of intervention A relative to intervention B as $\gamma_{AB} := \psi_A - \psi_B$. This formulation accommodates various types of outcomes, including binary, continuous, count, and censored event-time outcomes.

For example, in HIV prevention trials, the outcome of interest is often time to HIV acquisition, a censored event-time outcome. In such settings, the effect of intervention A relative to B is commonly summarized using the log hazard ratio (log HR), either assumed constant over time or evaluated at a fixed landmark time. Let $h_A(t)$ and $h_B(t)$ denote the hazard functions under interventions A and B , respectively, where $h_A(t)$ represents the instantaneous risk of HIV acquisition at time t for individuals receiving intervention A . At a fixed time t_0 , e.g., two years, the effect contrast γ_{AB} , i.e., the log hazard ratio, can be written as $\psi_A - \psi_B$ where $\psi_A := \log\{h_A(t_0)\}$. While our framework focuses on effect measures expressible as $\gamma_{AB} = \psi_A - \psi_B$, a commonly reported quantity in HIV prevention trials is prevention efficacy (PE), often defined as $1 - \text{HR} = 1 - \exp(\gamma_{AB})$. Although PE is not itself in the additive form we target, it is a monotonic transformation of the log HR and is thus implicitly included in analyses based on this contrast.

We denote the placebo, active control, and experimental intervention as P , C , and X respectively. Let γ_{XP} , γ_{CP} , and γ_{XC} represent their relative effects in the target population, i.e., in the population from which the active-controlled trial is sampled. The effect of the experimental intervention relative to placebo can be decomposed as $\gamma_{XP} = \gamma_{XC} + \gamma_{CP}$. When

negative values of the effect indicate a benefit, the scientific null hypothesis is:

$$H_0 : \quad \gamma_{XC} + \gamma_{CP} \geq \Delta, \quad (2.1)$$

where Δ denotes the null efficacy that must be ruled out. For the traditional preservation of effect non-inferiority criterion, $\Delta = f\gamma_{CP}$, with $f \in (0, 1)$ denoting the fraction of the active control effect to be preserved. For the novel inferred efficacy criterion, $\Delta = \Delta_0$, where Δ_0 is a fixed minimally acceptable level of efficacy. Under an inferred efficacy criterion, the goal is to determine if the experimental intervention meets the success criterion that would have been used, had a placebo-controlled design been possible.

The scientific null hypothesis (2.1) hinges on γ_{CP} , a parameter unidentifiable from the active-controlled trial alone. However, historical placebo-controlled trials provide evidence as to the active control effect in the historical setting. Let $\gamma_{CP,H}$ denote the active control's effect in the historical setting, reflecting its performance in the population from which the historical trials were drawn, administered following those trials' protocols, including adherence and relevant modifiers. The traditional synthesis method assumes $\gamma_{CP} = \gamma_{CP,H}$, a premise known as the constancy assumption (Wang et al., 2002; Hung et al., 2003; Fleming, 2008; FDA et al., 2016). The fixed-margin method assumes $\gamma_{CP} = \gamma_{CP,H} + a$, where $a \geq 0$ reflects a conservative adjustment based on the precision of the historical estimate and tends to zero as that estimate becomes more precise. Importantly, this adjustment is defined in the context of a conditional null hypothesis that depends on the observed historical data, whereas the null hypothesis we consider in this work is unconditional and accounts for uncertainty in the historical estimate. The specific form of a and this distinction will be detailed in Section 2.4.1. A standard assumption is also that normally distributed and consistent estimators $\hat{\gamma}_{XC}$ for γ_{XC} and $\hat{\gamma}_{CP,H}$ for $\gamma_{CP,H}$ exist with variances V_{XC} and $V_{CP,H}$, and that $\hat{\gamma}_{XC}$ and $\hat{\gamma}_{CP,H}$ are

independent. While the assumptions regarding the estimators are justifiable, being based on independent data from large-scale trials, the validity of the constancy assumption is questionable, and requires careful scrutiny. In the context of HIV prevention, disparities may arise due to variations in participants' adherence or in the susceptibility of circulating HIV viruses to the active control intervention.

Returning to the HIV prevention example, a reasonable estimate for the effect of CAB-LA among sub-Saharan African women is provided by [Donnell et al. \(2023\)](#), who used data from three contemporary studies conducted across five countries (Botswana, Kenya, Malawi, South Africa, and Zimbabwe). Their analysis adjusted for potential effect modifiers such as age and baseline diagnosis of sexually transmitted infections, specifically gonorrhea or chlamydia, to account for differences in sexual behavior. They estimated a log HR of $\hat{\gamma}_{CP,H} = \log(0.072) \approx -2.64$ (corresponding to a 92.8% prevention efficacy), with an estimated standard error of $\sqrt{V_{CP,H}} = 0.61$. Assuming constancy and setting $\gamma_{CP} = \hat{\gamma}_{CP,H}$, preserving 50% of the active control effect yields a null-efficacy threshold of $\Delta = -1.32$. Under the inferred efficacy criterion corresponding to 30% prevention efficacy, the threshold becomes $\Delta = \log(0.7) \approx -0.36$. This illustrates how different non-inferiority criteria translate into substantially different benchmarks for success.

2.4 The General Framework

This section introduces our general framework for evaluating non-inferiority methods. We begin by defining key parameters that allow us to translate a range of existing methods into this unified structure. The interpretations of these parameters are made relative to the scientific null hypothesis of interest considered in this framework, which may differ from the original motivations or hypotheses underlying those methods. We clarify these differences where relevant.

To quantify the degree of non-constancy that a given method can accommodate, we parameterize deviations from constancy. Let $\lambda_0 := (\gamma_{CP} - \gamma_{CP,H}) / \gamma_{CP,H}$ denote the *true relative effect deviation*. Here, $\lambda_0 = 0$ corresponds to constancy. Assuming that negative effect values indicate benefit, $\lambda_0 < 0$ means that the active control is less effective in the target population than in the historical population, while $\lambda_0 > 0$ indicates greater effectiveness. This parameterization yields $\gamma_{CP} = (1 + \lambda_0)\gamma_{CP,H}$, allowing the scientific null hypothesis (2.1) to be written as

$$H_0 : \quad \gamma_{XC} + (1 - f)(1 + \lambda_0)\gamma_{CP,H} \geq \Delta_0.$$

While this formulation replaces the unidentifiable parameter γ_{CP} with the relative effect deviation λ_0 , the latter remains unidentifiable. To address this, we introduce the parameter λ_1 , referred to as the *assumed relative effect deviation*. This assumption may be made explicitly—for example, by fixing $\lambda_1 = 0$ or assigning it a negative value—or implicitly through the way historical data are incorporated into the analysis.

Using λ_1 , we define the *operational null hypothesis* as:

$$H_0^{op} : \quad \gamma_{XC} + (1 - f)(1 + \lambda_1)\gamma_{CP,H} \geq \Delta_0. \tag{2.2}$$

This is the hypothesis that is actually tested using data from both the active-controlled trial and historical trials, as opposed to the scientific null hypothesis (2.1), which relies on unidentifiable parameters. Distinguishing between scientific and operational null hypotheses is important for comparing type-I error rates under non-constancy across methods, as these rates are evaluated at the boundary of the scientific null hypothesis.

Although one might object that we are simply replacing one unidentifiable parameter with another, working with λ_1 offers practical advantages. It provides a scale-free, interpretable way to characterize departures from constancy across different settings and clarifies how

methods incorporate historical data. Our framework makes these assumptions explicit and place them on a common scale. This will become clearer in the next subsection.

Within this general framework, we assess the operational null hypothesis (2.2) using the following test statistic:

$$T_{v,w,\lambda_1,f,\Delta_0} = \frac{\hat{\gamma}_{XC} + (1-f)(1+\lambda_1)\hat{\gamma}_{CP,H} - \Delta_0}{\sqrt{\widetilde{V}}}, \quad (2.3)$$

where

$$\widetilde{V} = V_{XC} + 2vw(1-f)(1+\lambda_1)\sqrt{V_{XC}V_{CP,H}} + w^2(1-f)^2(1+\lambda_1)^2V_{CP,H}.$$

Here, $v \in [0, 1]$ is the *unifying* parameter, accommodating both fixed-margin ($v = 1$) and synthesis methods ($v = 0$); and $w \geq 0$ serves as a *variance inflation* parameter, allowing variations and extensions of these methods. As discussed, λ_1 encapsulates our assumption about the relative effect deviation λ_0 , while f and Δ_0 specify the success criterion. In particular, $\Delta_0 = 0$ and $f \in (0, 1)$ yields the preservation of effect criterion, whereas $f = 0$ and $\Delta_0 < 0$ corresponds to the new inferred efficacy criterion.

In Appendix A.1, we show that at the boundary of the operational null hypothesis, i.e., when $\gamma_{XC} + (1-f)(1+\lambda_1)\gamma_{CP,H} = \Delta_0$, the test statistic $T_{v,w,\lambda_1,f,\Delta_0}$ is normally distributed generally with mean zero and variance $\sigma^2 = [V_{XC} + (1-f)^2(1+\lambda_1)^2V_{CP,H}] / \widetilde{V}$. We reject the operational null hypothesis when $T_{v,w,\lambda_1,f,\Delta_0}$ falls below $-Z_{1-\alpha}$, where $Z_{1-\alpha}$ denotes the $(1-\alpha)$ quantile of the standard normal distribution.

2.4.1 Special Cases within Our Framework

We now show that our general framework accommodates most existing methods for assessing non-inferiority (see Table 2.1).

Table 2.1: Common methods for assessing non-inferiority and their robustness to type I error control under constancy

Method	Parameter			UTIE control*
	Unifying (v)	Variance inflation (w)	Relative effect deviation (λ_1)	
Traditional synthesis method ¹		1	0	Yes
Bias-adjusted synthesis method ^{2,3}	0	1	λ_1	Yes†
Odem-Davis method ³		$1/(1 + \lambda_1)$	λ_1	Yes†
Fixed-margin 95-95 method ¹	1	1	0	Yes
Fixed-margin 0-95 method ¹		0	0	No

UTIE, unconditional type-I error. *Control indicates that the UTIE does not exceed the significant level α . For the 95-95 and 0-95 methods, $\alpha = 0.025$. †This condition is satisfied when $\lambda_1 < 0$, which is comon for these methods. References: 1. Wang et al. (2002); 2. Rothmann et al. (2003); Odem-Davis and Fleming (2013); and 3. Odem-Davis and Fleming (2013)

Synthesis Methods

Synthesis methods represent a simple strategy to incorporating historical information. These methods directly combine data from the historical and active-control trials, and importantly, their operational null hypothesis aligns exactly with that of (2.2).

The so-called *bias-adjusted synthesis method* assumes a relative effect deviation of λ_1 and directly combines $\hat{\gamma}_{XC}$ and $\hat{\gamma}_{CP,H}$ to construct an estimator of $\gamma_{XC} + (1 - f)(1 + \lambda_1)\gamma_{CP,H}$, with corresponding variance $V_{XC} + (1 - f)^2(1 + \lambda_1)^2V_{CP,H}$. The operational null hypothesis (2.2) is rejected if the upper bound of the $(1 - 2\alpha)\%$ CI for this quantity lies below Δ_0 . Specifically, this method rejects (2.2) when

$$\hat{\gamma}_{XC} + (1 - f)(1 + \lambda_1)\hat{\gamma}_{CP,H} + Z_{1-\alpha}\sqrt{V_{XC} + (1 - f)^2(1 + \lambda_1)^2V_{CP,H}} < \Delta_0,$$

which is equivalent to the following test statistic

$$T_{SM,\lambda_1} = \frac{\hat{\gamma}_{XC} + (1 - f)(1 + \lambda_1)\hat{\gamma}_{CP,H} - \Delta_0}{\sqrt{V_{XC} + (1 - f)^2(1 + \lambda_1)^2V_{CP,H}}}$$

being less than $-Z_{1-\alpha}$. This corresponds to our general test statistic in (2.3) with parameters $(v, w, \lambda_1) = (0, 1, \lambda_1)$. When constancy is assumed—i.e., when $\lambda_1 = 0$ —we refer to this as the *traditional synthesis method*.

Given that synthesis methods lack robustness to deviations beyond the assumed relative effect deviation λ_1 —particularly in their control of type I error (as will be clarified in the next section)—Odem-Davis and Fleming (2013) proposed a modification that avoids contracting the variance $(1-f)\hat{\gamma}_{CP,H}$ by the factor $(1+\lambda_1)^2$. This variant uses the same point estimator for $\gamma_{XC} + (1-f)(1+\lambda_1)\gamma_{CP,H}$, but assumes a larger variance $V_{XC} + (1-f)^2V_{CP,H}$. The resulting test statistic is

$$T_{OD,\lambda_1} = \frac{\hat{\gamma}_{XC} + (1-f)(1+\lambda_1)\hat{\gamma}_{CP,H} - \Delta_0}{\sqrt{V_{XC} + (1-f)^2V_{CP,H}}},$$

which corresponds to the general form (2.3) with parameters $(v, w, \lambda_1) = (0, 1/(1+\lambda_1), \lambda_1)$. Notably, this method is only applicable and meaningful when $\lambda_1 \neq 0$. We refer to it as the *Odem-Davis method*.

Fixed Margin Methods

Fixed-margin methods are the most widely used methods for incorporating historical information in non-inferiority testing (Wang et al., 2002; Fleming, 2008; Rothmann et al., 2011; FDA et al., 2016). Unlike synthesis methods, a fixed-margin method treats the historical data as fixed and known, and assumes that γ_{CP} can be conservatively approximated by $\tilde{\gamma}_{CP} := \hat{\gamma}_{CP} + a$ where $a = Z_{1-\theta}\sqrt{V_{CP,H}}$ and $\theta \in (0, 0.5]$, typically $\theta = \alpha$. The term a represents the absolute deviation from the point estimate $\hat{\gamma}_{CP,H}$ used to construct the upper bound of a $100(1-2\theta)\%$ CI for $\gamma_{CP,H}$. This serves as a conservative adjustment to account for uncertainty in the historical estimate. A fixed success margin is then defined as

$\delta := \Delta_0 - (1-f)\tilde{\gamma}_{CP}$. The method tests the following conditional operational null hypothesis:

$$H_0^{\text{FM}} : \quad \gamma_{XC} \geq \delta. \quad (2.4)$$

This null hypothesis is defined conditional on the historical data and corresponds to rejecting the scientific null (2.1) under the assumption that $\gamma_{CP} \leq \tilde{\gamma}_{CP}$ —i.e., that any possibly downward deviation from constancy has already been absorbed into the conservative margin δ .

A fixed-margin method rejects (2.4) if the upper bound of the $100(1 - 2\alpha)\%$ CI for γ_{XC} falls below δ , i.e., if $\hat{\gamma}_{XC} + Z_{1-\alpha}\sqrt{V_{XC}} < \Delta_0 - (1-f) [\hat{\gamma}_{CP,H} + Z_{1-\theta}\sqrt{V_{CP,H}}]$. This is equivalent to the test statistic

$$T_{FM,\theta,\lambda_1} = \frac{\hat{\gamma}_{XC} + (1-f)\hat{\gamma}_{CP,H} - \Delta_0}{\sqrt{V_{XC}} + \frac{Z_{1-\theta}}{Z_{1-\alpha}}(1-f)\sqrt{V_{CP,H}}}.$$

being less than $-Z_{1-\alpha}$. Within our framework, this corresponds to setting $(v, w, \lambda_1) = (1, Z_{1-\theta}/Z_{1-\alpha}, 0)$ in the general statistic (2.3). When $\theta = \alpha = 0.025$, this yields the well-known *95-95 method*. When $\theta = 0.5$ and $\alpha = 0.025$, it corresponds to the *0-95 method*, which assumes no effect deviation from the historical estimate.

An important subtlety arises when placing fixed-margin methods within our framework. While these methods coherently define and test a conditional null hypothesis based on the observed historical data, their translation into our unconditional framework may initially seem inconsistent—but this is not the case. The conservative estimate $\tilde{\gamma}_{CP}$ used to define δ reflects an implicit belief that γ_{CP} may fall below its historical estimate. In our framework, this is captured by assuming no effect deviation, i.e., $\lambda_1 = 0$, and inflating the variance to account for uncertainty in the historical estimate. In this sense, fixed-margin methods

resemble synthesis methods with $\lambda_1 = 0$ but with an adjusted variance, scaled by:

$$\frac{V_{XC} + 2(1-f)\frac{Z_{1-\theta}}{Z_{1-\alpha}}\sqrt{V_{XC}V_{CP,H}} + \left(\frac{Z_{1-\theta}}{Z_{1-\alpha}}\right)^2(1-f)^2V_{CP,H}}{V_{XC} + \left(\frac{Z_{1-\theta}}{Z_{1-\alpha}}\right)^2(1-f)^2V_{CP,H}},$$

which depends on both the uncertainty in the historical estimate of $\gamma_{CP,H}$ and that of γ_{XC} in the active-controlled trial. Because this scaling depends on an unknown quantity, fixed-margin methods are not formally considered synthesis methods within our classification.

The Unified Framework of Snapinn and Jiang (2008)

Snapinn and Jiang (2008) proposed a unified class of methods that encompasses many existing non-inferiority procedures for the specific success criterion defined by $\Delta_0 = 0$ and $f = 0$. This class is fully nested within our general framework when the parameters are set to $(w, f, \Delta_0) \equiv (1, 0, 0)$. However, even within this restricted setting, their class does not accommodate certain procedures—most notably, the Odem–Davis method—which fall outside the structure they considered.

2.5 Operating Characteristics within the General Framework

Two types of operating characteristics are pertinent for evaluating non-inferiority in active-controlled trials: conditional and unconditional. Both account for stochasticity in the active-controlled trial, but conditional characteristics assume a fixed and known active control effect in the historical population, while unconditional ones treat $\hat{\gamma}_{CP,H}$ as a random variable with a known variance $V_{CP,H}$. Since unconditional operating characteristics incorporate uncertainty in estimating the active control effect, robustness to constancy assumption violations is most naturally evaluated in terms of unconditional type-I error. We evaluate unconditional operating characteristics and facilitate designing an active-controlled trial to achieve adequate

unconditional power while controlling unconditional type-I error.

Under the assumptions outlined in the preceding sections, the unconditional power of the statistic (2.3) to reject the scientific null hypothesis (2.1) when the experimental intervention has an effect size γ_{XP} is

$$\Phi \left(\frac{\Delta_0 + \{(1 + \lambda_0) - (1 - f)(1 + \lambda_1)\} \gamma_{CP,H} - \gamma_{XP} - Z_{1-\alpha} \sqrt{\widetilde{V}}}{\sqrt{V_{XC} + (1 - f)^2 (1 + \lambda_1)^2 V_{CP,H}}} \right), \quad (2.5)$$

where Φ is the cumulative distribution function of the standard normal distribution. Its unconditional type-I error at the boundary of the scientific null hypothesis (2.1) is

$$\Phi \left(\frac{(1 - f)(\lambda_0 - \lambda_1) \gamma_{CP,H} - Z_{1-\alpha} \sqrt{\widetilde{V}}}{\sqrt{V_{XC} + (1 - f)^2 (1 + \lambda_1)^2 V_{CP,H}}} \right). \quad (2.6)$$

Similarly, the conditional power and type-I error are

$$\Phi \left(\frac{\Delta_0 + \{(1 + \lambda_0) - (1 - f)(1 + \lambda_1)\} \hat{\gamma}_{CP,H} - \gamma_{XP} - Z_{1-\alpha} \sqrt{\widetilde{V}}}{\sqrt{V_{XC}}} \right) \quad (2.7)$$

and

$$\Phi \left(\frac{(1 - f)(\lambda_0 - \lambda_1) \hat{\gamma}_{CP,H} - Z_{1-\alpha} \sqrt{\widetilde{V}}}{\sqrt{V_{XC}}} \right). \quad (2.8)$$

For completeness, we provide the derivations of expressions (2.5)–(2.8) in Appendix A.1.

To explore the implications of these expressions in practical settings, we focus on a specific region of the parameter space where power is at least 50% and type-I error is below 50%. In Appendix A.2, we show that the parameter space of interest satisfies

$$\gamma_{XP} < \Delta_0 + (1 + \lambda_0) \gamma_{CP,H} - (1 - f)(1 + \lambda_1) \left\{ \gamma_{CP,H} + w Z_{1-\alpha} \sqrt{V_{CP,H}} \right\} \quad (2.9)$$

and

$$\lambda_1 \leq \lambda_0 - \frac{(1 + \lambda_0)wZ_{1-\alpha}\sqrt{V_{CP,H}}}{\gamma_{CP,H} + wZ_{1-\alpha}\sqrt{V_{CP,H}}}. \quad (2.10)$$

We further assume that the active control is efficacious in the historical population, i.e., $\gamma_{CP,H} < 0$.

Under these conditions, conditional power always exceeds unconditional power, and conditional type-I error is always lower than unconditional type-I error. This occurs because conditional characteristics ignore the uncertainty associated with estimating the active control effect from historical data. Mathematically, this follows from two key observations: first, expressions (2.5) and (2.6) reduce to (2.7) and (2.8) when $V_{CP,H} = 0$ and $\hat{\gamma}_{CP,H} = \gamma_{CP,H}$; and second, under constraints (2.9) and (2.10), the numerators inside the standard normal arguments of (2.5) and (2.6) are strictly positive and strictly negative, respectively.

Moreover, both power and type-I error decrease as λ_0 increases. Mathematically, this inverse relationship arises because the numerators inside the standard normal arguments of (2.5) and (2.6) depend inversely on λ_0 , given that $\gamma_{CP,H}$ is assumed negative. Clinically, this makes sense: an increase from λ_0 to λ'_0 implies either (i) less deviation from constancy if $\lambda'_0 < 0$, or (ii) that the active control is more effective in the target population than in the historical one if $\lambda'_0 \geq 0$. In the first case, the historical estimate becomes more accurate, and in the second, it serves as a conservative approximation of the active control effect in the target population. In both cases, this leads to reduced type-I error. However, it also reduces power, since the null hypothesis becomes harder to reject when the active control is more effective in the target population.

In the following section, we derive additional insights from expressions (2.5)–(2.8) and introduce further operating characteristics relevant to trial design—both conceptually and through an application of the proposed framework to the design of an HIV prevention trial.

2.6 Application to Designing an HIV Prevention Trial

As introduced in Section 2.2, the goal is to design an HIV prevention trial in sub-Saharan African women using CAB-LA as the active control. The outcome of interest is time to HIV acquisition, with the log hazard ratio used as the measure of effect throughout our methodological framework and for computing all operating characteristics. However, because results on the log HR scale are difficult to interpret clinically, we present all relevant quantities on the prevention efficacy scale. Given that CAB-LA has demonstrated high efficacy in the historical population (PE of 92.8%, 95% CI: 76.1%–97.8%), we consider both type of non-inferiority success criterion: i) the 50% preservation of effect criterion, parameterized in our framework as $(f, \Delta_0) = (0.5, 0)$, and ii) the inferred 30% prevention efficacy criterion, parameterized as $(f, \Delta_0) = (0, \log(0.7))$ on the log HR scale.

We explore five analytical methods: i) the traditional synthesis method, ii) the bias-adjusted synthesis method with $\lambda_1 = -0.23$, iii) the Odem-Davis method with $\lambda_1 = -0.23$, iv) the 95-95 method, v) and the 0-95 method. The value $\lambda_1 = -0.23$ used in the bias-adjusted synthesis and Odem-Davis methods corresponds to a CAB-LA PE of 86.8%, which is one standard deviation below the estimated active control effect on the log HR scale.

2.6.1 Maximum Unconditional Power

Incorporating the uncertainty in estimating the active control effect from historical data reduces power. For a design alternative within the restricted parameter space specified in Section 2.5, a study can be carried out to achieve any conditional power above 50%. However, there is an upper bound to the corresponding unconditional power, given by

$$\Phi \left(-wZ_{1-\alpha} + \frac{\Delta_0 + \{(1 + \lambda_0) - (1 - f)(1 + \lambda_1)\} \gamma_{CP,H} - \gamma_{XP}}{(1 - f)(1 + \lambda_1) \sqrt{V_{CP,H}}} \right).$$

Details of this derivation are shown in Appendix A.2. This means that some design alternatives cannot be detected unconditionally with the desired power, affecting the interpretation of findings from active-controlled trials designed to control conditional power.

Figure 2.1 shows the maximum unconditional power for detecting interventions with prevention efficacy between 65% and 98% across the five analytical methods. We see, for example, a 90% PE intervention cannot be detected with 90% unconditional power with most methods, except with the 0-95 method, while a 95% PE intervention can be detected with 90% unconditional power across all methods and criteria. This underscores the challenges of non-inferiority designs when the active control is highly effective.




Figure 2.1: Maximum unconditional power for detecting design alternatives across our selected methods and criteria. SM, synthesis method; BA-SM, bias-adjusted synthesis method; OD, Odem-Davis method. The horizontal dashed line represents an unconditional power of 0.90. The vertical dotted line represents an experimental intervention with a PE of 95%.

2.6.2 Trial Design Controlling Unconditional Power

Active-controlled trials for assessing non-inferiority are typically designed to achieve a target conditional power. However, as previously discussed, conditional power is systematically lower than unconditional power. This discrepancy may result in trials being underpowered, leading to the failure to approve interventions that are in fact non-inferior. We refer to this conventional practice as the *traditional approach* to clinical trial design.

To mitigate the loss of power due to potential deviations from the constancy assumption (i.e., increases in λ_0), a common ad hoc strategy is to assume that the active control is more effective in the target population than suggested by historical data. This assumption increases the conditional power during planning and is referred to here as the *ad hoc approach*.

While calculating the sample size to achieve a desired level of conditional power is a well-established procedure (see, e.g., Rothmann et al. (2011)), our framework also enables the design of trials that control unconditional power—the probability of rejecting the scientific null hypothesis when both historical and trial data are treated as random, assuming a fixed deviation from constancy, i.e., a fixed value of λ_0 . We refer to this as the *novel approach*.

Designing a non-inferiority trial to ensure a target level of unconditional power to detect a given design alternative is a multi-step approach. First, determine if the desired power for a design alternative can be achieved through conditions in the preceding section. Second, if achievable, calculate the numerical value for V_{XC} by solving formula (2.5) equated to the desired power. Third, determine the sample size N based on the relationship between V_{XC} and sample size in the active-controlled trial. Fourth, evaluate the accommodated level of non-constancy using expression (2.12) as described in subsection 2.6.4.

We apply this procedure to evaluate three trial design strategies: 1) the traditional approach, which ensures 90% conditional power, 2) the novel approach, which guarantees 90% unconditional power, and 3) the ad hoc approach that controls conditional power under

the assumption that CAB-LA has 94.7% PE—an effect half a standard deviation above the estimated active control effect on the log scale. Our design alternative is an experimental intervention with 95% PE, which we confirmed to be feasible in the previous subsection. The required number of events is computed under constancy, i.e., $\lambda_0 = 0$, and 1:1 allocation. Total sample size is then determined assuming a fixed two-year follow-up period, a 3% baseline incidence under placebo, 7.5% annual loss to follow-up, and constant incidence and PE in both groups.

The trial design specifications and operating characteristics for the various approaches, methods, and success criteria considered in the HIV prevention trial are summarized in Table 2.2 for the preservation of effect criterion and in Table 2.3 for the inferred efficacy criterion. To facilitate interpretation, we first introduce two additional operating characteristics within our general framework: the success margin and the tolerable level of non-constancy. We also provide guidance on comparing different success criteria, which is important to consider before reviewing the results presented in the tables.

2.6.3 Establishing a Success Margin

Our framework elucidates that any method within its scope can be regarded as a fixed-margin method. From the definition of our general test statistic (2.3), we find that the rejection region $T_{v,c,f,\Delta_0} < -Z_{1-\alpha}$ is equivalent to $\hat{\gamma}_{XC} + Z_{1-\alpha}\sqrt{V_{XC}} < \delta$, where

$$\delta = \Delta_0 - (1 - f)(1 + \lambda_1)\hat{\gamma}_{CP,H} - Z_{1-\alpha}\left\{\sqrt{\tilde{V}} - \sqrt{V_{XC}}\right\}. \quad (2.11)$$

Notably, in general, the margin depends on V_{XC} , and thus relies on data from the active-controlled trial. However, for fixed margin methods ($v = 1$), the margin simplifies to $\Delta_0 - (1 - f)(1 + \lambda_1)\left\{\hat{\gamma}_{CP,H} + wZ_{1-\alpha}\sqrt{V_{CP,H}}\right\}$, which only depends on historical data and

design parameters. Having a defined form for the success margin enables the investigator to characterize at the design stage the size of the experimental to active control effect that would constitute success, for potential values of V_{XC} .

The success margins in our HIV prevention trial design example are reported in the second column of Tables 2.2 and 2.3, expressed on the relative hazard ratio scale. These margins represent the upper bound for the relative hazard ratio between the experimental and active control interventions, below which the experimental intervention is deemed successful.

2.6.4 Quantifying Controlled Non-Constancy

Given the need to ensure strict control over type-I error in phase 2b/3 trials and the result that non-constancy can inflate this error, quantifying the level of non-constancy accommodated is crucial. As the true extent of non-constancy remains unknown, our best course of action is to establish thresholds on the maximum degree of non-constancy—or, in simpler terms, the minimum value of λ_0 —where the unconditional type-I error remains at or below α . We call this the *tolerable level of non-constancy* which is shown in Appendix A.3 to be:

$$\lambda_{0,\min} = \lambda_1 + \frac{\sqrt{\widetilde{V}} - \sqrt{V_{XC} + (1-f)^2(1+\lambda_1)^2V_{CP,H}}}{(1-f)\gamma_{CP,H}} Z_{1-\alpha}. \quad (2.12)$$

The tolerable level of non-constancy allows practitioners to evaluate and compare the robustness against non-constancy of different designs and success criteria (see Table 2.1). The traditional synthesis method has $\lambda_{0,\min} = 0$, ensuring protected type-I error under constancy but inflated error when $\lambda_0 < 0$. The 95-95 method has $\lambda_{0,\min}$ strictly smaller than 0. Hence, it has protected type-I error under constancy, and allows for some degree of non-constancy. The amount of non-constancy accommodated depends on the uncertainty in $\hat{\gamma}_{CP,H}$ as shown in Appendix A.3. The 0-95 method has $\lambda_{0,\min} > 0$, making it anti-conservative under con-

stancy. The bias-adjusted synthesis method has $\lambda_{0,\min} = \lambda_1$; the Odem-Davis method has $\lambda_{0,\min}$ strictly smaller than λ_1 . Hence, both methods have protected (conservative) type-I error when the non-constancy level is less extreme than the assumed non-constancy level λ_1 . Notably, $\lambda_1 < 0$ means that the active control in the target population is assumed to be less effective than in the historical population.

The tolerable levels of non-constancy for the different methods examined in our HIV prevention trial design example are provided in the CNC—controlled non-constancy—column of Tables 2.2 and 2.3, expressed on the prevention efficacy scale. These values represent the lowest true prevention efficacy of CAB-LA in the target population for which unconditional type I error is controlled when rejection is based on the success margin.

2.6.5 Comparing Success Criteria

A final insight emerges from the general framework. While the formulas (2.5)–(2.8) enable numerical comparisons of operating characteristics across different success criteria, it is essential to ensure that their scientific null hypotheses are aligned numerically for a fair comparison in terms of type-I error. Without such alignment, the type-I error rates are not comparable, as they correspond to different scientific null hypotheses.

For the preservation of effect criterion, the scientific null hypothesis depends on the true active control effect within the target population. Thus, even if a preservation of effect criterion and an inferred efficacy criterion are aligned so that they correspond to the same scientific null hypothesis, this alignment only holds for a particular value of the active control effect. Consequently, we advocate for selecting the success criterion based on scientific rationale rather than solely based on a comparison of operating characteristics.

In our HIV prevention trial design example, this consideration informed our use of two distinct criteria. Preserving 50% of a 92.8% CAB-LA PE on the log scale corresponds

to a 73.2% inferred PE, while requiring an inferred PE of 30% preserves only 13.6% of the CAB-LA effect. These targets are not numerically aligned, but they reflect different scientific goals and policy questions—one prioritizing relative preservation of an established intervention’s efficacy, the other focusing on achieving a minimum level of absolute protection with meaningful public health impact.

2.6.6 Systematic Comparison within the General Framework

We now walk through the trial design results in Tables 2.2 and 2.3, highlighting important comparisons. For example, using the 95-95 method with the 50% preservation of effect criterion, achieving 90% unconditional power requires 15,516 participants (7,758 per arm) to accrue 52 events. The experimental intervention is deemed successful if the hazard of HIV acquisition in this group is up to 2.05 times that in the control group. This design controls unconditional type-I error below 0.025 as long as CAB-LA PE is at least 85.9% ($\lambda_{0,\min} = -0.26$), and it reaches 70% unconditional power when the true CAB-LA PE is higher than expected, at 94.7%. In contrast, the conditional approach requires 11,012 but achieves only 83% unconditional power under constancy and 60% unconditional power if CAB-LA PE is 94.7%.

Evaluating these trial designs shows that sizing a trial to achieve 90% conditional power yields unconditional power below 90%, confirming the risk of underpowered trials. The ad hoc approach, targeting 90% conditional power under a higher CAB-LA PE of 94.7%, generally requires sample sizes 1.5 times larger than the traditional approach based on conditional power, often achieving over 90% unconditional power but at a higher resource cost. The approach controlling unconditional power strikes a balance between the two aforementioned approaches. Additionally, the unconditional approach demonstrates greater robustness to non-constancy in terms of unconditional power compared to the traditional approach. Un-

conditional power consistently exceeds 70% using the preservation of effect criterion, and exceeds 75% with the inferred efficacy criterion, when CAB-LA PE is 94.7%. This underscores its utility in ensuring adequate power while maintaining manageable sample sizes and ensuring robustness to non-constancy across various scenarios.

When comparing the five methods within the same success criterion and design approach, the 0-95 method yields the most lenient margin and the smallest sample size, while demonstrating the greatest robustness to non-constancy in terms of loss of power. However, as anticipated, it inflates type-I error even under constancy ($\lambda_{0,\min} > 0$). In contrast, the Odem-Davis method accommodates the highest level of non-constancy while maintaining proper type-I error control.

Lastly, for a given method and design approach, the 50% preservation of effect criterion sets stricter margins and typically requires larger sample sizes compared to the inferred 30% efficacy criterion, except with the Odem-Davis method under unconditional power control.

Overall, this application shows that: 1) with reasonable assumptions, the bias-adjusted synthesis method is efficient and accommodates substantial non-constancy, 2) the 95-95 method does not generally outperform others in robustness to non-constancy, 3) the inferred efficacy non-inferiority criterion is attractive for its interpretability and practicality, and 4) the approach that controls unconditional power effectively addresses uncertainty in the active control effect.

Table 2.2: Trial design specifications and operating characteristics for detecting 95% prevention efficacy with 90% power under the 50% effect preservation criterion, across methods and design strategies.

Traditional approach: Controlling conditional power and assuming CAB-LA PE = 92.8%						
Method	Margin*	RNE	Sample size†	CNC‡	UP0	UPa
Traditional SM	3.12	19 (8:11)	5,766 (2,883/arm)	92.8%	0.86	0.69
BA-SM, $\lambda_1 = -23\%$	2.42	27 (11:16)	8,008 (4,004/arm)	86.8%	0.86	0.65
OD with $\lambda_1 = -23\%$	2.21	32 (13:19)	9,510 (4,755/arm)	84.4%	0.86	0.63
95-95 method	2.05	37 (15:22)	11,012 (5,506/arm)	85.0%	0.83	0.60
0-95 method	3.73	15 (6:9)	4,504 (2,252/arm)	94.8%	0.87	0.72
Novel approach: Controlling unconditional power and assuming CAB-LA PE = 92.8%						
Method	Margin*	RNE	Sample size†	CNC‡	UP0	UPa
Traditional SM	3.07	24 (10:14)	7,208 (3,604/arm)	92.8%	0.90	0.75
BA-SM, $\lambda_1 = -23\%$	2.40	33 (14:19)	10,090 (5,045/arm)	86.8%	0.90	0.70
OD with $\lambda_1 = -23\%$	2.17	40 (16:24)	12,012 (6,006/arm)	84.3%	0.90	0.69
95-95 method	2.05	52 (21:31)	15,516 (7,758/arm)	85.9%	0.90	0.70
0-95 method	3.73	17 (7:10)	5,046 (2,523/arm)	94.9%	0.90	0.77
Ad hoc approach: Controlling conditional power and assuming CAB-LA PE = 94.7%						
Method	Margin*	RNE	Sample size†	CNC‡	UP0	UPa
Traditional SM	2.98	33 (16:17)	11,558 (5,779/arm)	92.8%	0.95	0.83
BA-SM, $\lambda_1 = -23\%$	2.32	53 (26:27)	18,738 (9,369/arm)	86.8%	0.97	0.91
OD with $\lambda_1 = -23\%$	2.06	71 (35:36)	25,226 (12,613/arm)	83.8%	0.97	0.82
95-95 method	2.05	72 (35:37)	25,226 (12,613/arm)	86.8%	0.95	0.78
0-95 method	3.73	23 (11:12)	8,160 (4,080/arm)	95.1%	0.95	0.85

SM, synthesis method; BA-SM, bias-adjusted synthesis method; OD, Odem-Davis method; RNE, required number of events assuming 1:1 randomization. The numbers in parentheses correspond to the expected number of events in each group. CNC, controlled non-constancy; UP0, unconditional power under constancy (when CAB-LA PE is 92.8%); UPa, unconditional power when CAB-LA PE is 94.7%. *The experimental intervention is deemed successful if the hazard ratio of HIV acquisition for the experimental intervention relative to the active control is lower than this margin. †Assuming 3% underlying placebo incidence, a fixed 2-year follow-up, 7.5% loss to follow-up per year, and constant incidence and PE. ‡This is the lowest CAB-LA PE such that the unconditional type-I error is below 0.025. It corresponds to $1 - \exp([1 + \lambda_{0,\min}]\hat{\gamma}_{CP,H})$.

Table 2.3: Trial design specifications and operating characteristics for detecting 95% prevention efficacy with 90% power under the 30% prevention efficacy criterion, across methods and design strategies.

Traditional approach: Controlling conditional power and assuming CAB-LA PE = 92.8%						
Method	Margin*	RNE	Sample size†	CNC‡	UP0	UPa
Traditional SM	6.13	9 (4:5)	2,882 (1,441/arm)	92.8%	0.83	0.73
BA-SM, $\lambda_1 = -23\%$	3.72	15 (6:9)	4,504 (2,252/arm)	86.8%	0.83	0.69
OD with $\lambda_1 = -23\%$	2.88	22 (9:13)	6,506 (3,253/arm)	83.7%	0.81	0.65
95-95 method	2.94	21 (9:12)	6,486 (3,243/arm)	87.0%	0.78	0.63
0-95 method	9.72	7 (3:4)	2,162 (1,081/arm)	95.2%	0.85	0.76
Novel approach: Controlling unconditional power and assuming CAB-LA PE = 92.8%						
Method	Margin*	RNE	Sample size†	CNC‡	UP0	UPa
Traditional SM	5.67	14 (6:8)	4,324 (2,162/arm)	92.8%	0.90	0.81
BA-SM, $\lambda_1 = -23\%$	3.51	24 (10:14)	7,208 (3,604/arm)	86.8%	0.90	0.78
OD with $\lambda_1 = -23\%$	2.50	49 (20:29)	14,514 (7,257/arm)	83.2%	0.90	0.76
95-95 method	2.94	48 (20:28)	14,414 (7,207/arm)	88.8%	0.90	0.79
0-95 method	9.72	8 (3:5)	2,502 (1,251/arm)	95.3%	0.90	0.83
Ad hoc approach: Controlling conditional power and assuming CAB-LA PE = 94.7%						
Method	Margin*	RNE	Sample size†	CNC‡	UP0	UPa
Traditional SM	5.72	13 (6:7)	4,760 (2,380/arm)	92.8%	0.89	0.81
BA-SM, $\lambda_1 = -23\%$	3.47	25 (12:13)	8,840 (4,420/arm)	86.8%	0.91	0.80
OD with $\lambda_1 = -23\%$	2.52	45 (22:23)	15,856 (7,928/arm)	83.2%	0.90	0.76
95-95 method	2.94	33 (16:17)	11,558 (5,779/arm)	88.1%	0.86	0.74
0-95 method	9.72	8 (4:4)	2,882 (1,441/arm)	95.4%	0.91	0.84

SM, synthesis method; BA-SM, bias-adjusted synthesis method; OD, Odeh-Davis method; RNE, required number of events assuming 1:1 randomization. The numbers in parentheses correspond to the expected number of events in each group. CNC, controlled non-constancy; UP0, unconditional power under constancy (when CAB-LA PE is 92.8%); UPa, unconditional power when CAB-LA PE is 94.7%. *The experimental intervention is deemed successful if the hazard ratio of HIV acquisition for the experimental intervention relative to the active control is lower than this margin. †Assuming 3% underlying placebo incidence, a fixed 2-year follow-up, 7.5% loss to follow-up per year, and constant incidence and PE. ‡This is the lowest CAB-LA PE such that the unconditional type-I error is below 0.025. It corresponds to $1 - \exp([1 + \lambda_{0,\min}]\hat{\gamma}_{CP,H})$.

2.7 Discussion

In this chapter, we present a general framework for the design of active-controlled trials evaluating the non-inferiority of an experimental intervention that encompasses a wide variety of methods and success criteria, and accommodates binary, continuous, count, and censored event-time outcomes. A key contribution of our framework is facilitating rational decision-making: it allows researchers to select the most suitable design based on the trade-off between type-I error and power and to evaluate the level of non-constancy that is accommodated by the design. The framework includes the novel inferred efficacy non-inferiority criterion as a special case. Additionally, it introduces a novel approach to trial design, prioritizing the control of unconditional power to ensure robustness to non-constancy.

The development of a framework for designing non-inferiority trials to evaluate the inferred efficacy criterion will advance future trial design. In particular, when an active control is highly efficacious, and an experimental intervention offers advantages in terms of acceptability or feasibility, a traditional preservation of effect margin may be overly stringent. The inferred efficacy criterion is also highly interpretable.

While our framework encompasses both the traditional preservation of effect criterion and the innovative inferred efficacy criterion, caution is warranted when directly comparing these criteria, as they typically correspond to different scientific null hypotheses. This distinction is important because previous studies suggest that preservation of effect criteria may demonstrate greater robustness to non-constancy violations than the inferred efficacy criterion (Wang et al., 2002). However, our work suggests that this is not true in general. Moreover, even if attempts are made to align the scientific null hypotheses, as shown in our illustrative example, this alignment is only valid for a specific value of the active control effect. This underscores the need to carefully consider the nuances of each criterion and their associated scientific null hypotheses when designing the trial and interpreting its outcomes.

Importantly, we considered the setting in which an active control is effective and available to the population of interest. However, in settings where the active control is not available or not acceptable to the population, a placebo-controlled design remains the gold standard for generating evidence regarding efficacy of the experimental intervention.

Overall, our framework advances the understanding and application of statistical methods for non-inferiority assessment, offering researchers a powerful toolkit for rigorous non-inferiority study design.

Chapter 3

PROXIMAL CAUSAL INFERENCE FOR MODIFIED TREATMENT POLICIES

3.1 *Introduction*

The methodological developments in this article are motivated by challenges arising in vaccine research. The development of vaccines requires rigorous assessment, including phase 3 efficacy randomized trials that often involve lengthy follow-up periods, delaying the rapid deployment of new, effective vaccines. To address this challenge, researchers aim to identify early vaccine-induced immune markers that predict later outcomes, known as correlates of protection (Plotkin, 2010; Fleming and Powers, 2012). These correlates can serve as surrogates in future studies, potentially expediting vaccine development and approval processes while maintaining safety and efficacy standards. Identifying correlates of protection involves characterizing the causal effect of a continuous exposure—the immune marker—on the outcome of interest. A natural approach to assessing this effect is through the causal dose-response curve, which assigns to each immune marker level the mean outcome that would be observed in a hypothetical scenario where everyone’s immune marker level is set to that value.

An alternative approach that has gained popularity in recent years involves the shift-intervention-response curve (Hejazi et al., 2021; Huang et al., 2023). This curve assigns to each possible shift the mean outcome that would be observed if everyone’s level of immune markers were set to their actual vaccine-induced value, adjusted by the specified shift. These interventions are of interest because for reasonable shift values, the modified immune marker

levels could realistically be achieved with alternative vaccine formulations.

A shift intervention is an example of a modified treatment policy (MTP) (Robins, 1986; Díaz Muñoz and Van Der Laan, 2012; Haneuse and Rotnitzky, 2013). Quite generally, an MTP is an intervention that assigns each individual a treatment dose corresponding to their observed dose in the study but modified in a pre-specified manner. Importantly, MTPs allow treatment assignment to depend on baseline covariates, enabling more flexible and personalized interventions. For example, the amount of dose shift could vary based on individual characteristics such as age, immunocompetency status, or other relevant factors.

An additional complication when assessing immune markers as correlates of protection is that vaccine-induced immune marker levels cannot be randomized, making any analysis with immune marker levels as the exposure susceptible to confounding, possibly due to unmeasured confounders. In particular, an individual’s underlying immunity level is likely a strong unobserved confounder. A promising solution to this challenge is the proximal causal inference framework, recently introduced by Miao et al. (2018); Tchetgen et al. (2020); and Cui et al. (2024). This framework enables the identification and inference of causal effects in scenarios involving unobserved confounders, provided one has access to two disjoint sets of observed strong proxies for the unmeasured confounders. These proxies are commonly referred to as negative control treatments and negative control outcomes. Negative control treatments are proxies of the unmeasured confounders that lack a direct impact on the outcome of interest, while negative control outcomes are proxies that remain unaffected by both the negative control treatments and the primary treatment of interest. Figure 3.1 provides a visual representation of a causal diagram incorporating these negative control variables, illustrating their relationships within the framework. The proximal causal inference framework is well-suited for identifying correlates of protection in vaccine efficacy studies, as these trials routinely measure variables that qualify as negative controls, making the

framework’s application both feasible and relevant. While estimators of the counterfactual MTP mean under the assumption of no unmeasured confounders have been extensively studied (Díaz Muñoz and Van Der Laan, 2012; Haneuse and Rotnitzky, 2013; Hejazi et al., 2021), to our knowledge, estimators under the proximal causal inference assumptions have not yet been developed for this context. In this chapter, we address this gap by proposing a novel estimator of the counterfactual MTP mean that leverages the proximal causal inference framework to assess immune markers as correlates of protection in vaccine efficacy studies.

In Section 3.2, we present our motivating example, the ENSEMBLE trial, highlighting the proxies used to address unmeasured confounding. Section 3.3 introduces the causal estimand of interest and the assumptions that allow for its identification. In Section 3.4, we describe our estimator and provide its asymptotic analysis. Section 3.5 evaluates the finite sample properties of our estimator through numerical experiments. In Section 3.6, we apply our estimator to identify correlates of protection in the ENSEMBLE trial. Finally, we offer general conclusions and future research directions in Section 3.7.




Figure 3.1: A causal directed acyclic graph of proximal causal inference. A and Y represent the exposure and outcome of interest; L and U the observed and unobserved confounders; and Z and W the negative control treatments and negative control outcomes, respectively.

3.2 Motivating Example

We aim to investigate vaccine-induced neutralizing antibody levels at day 29 (D29) post-vaccination as correlates of protection against symptomatic COVID-19 infection in the first 7 months post antibody measurement. This immune marker has been identified as a correlate of protection across various vaccine platforms through multiple research approaches (Feng et al., 2021; Gilbert et al., 2022a,b; Fong et al., 2022, 2023). Our analysis utilizes data from the ENSEMBLE trial, a multinational, randomized, placebo-controlled phase 3 study assessing the efficacy of a single dose of Ad26.COV2.S in preventing symptomatic SARS-CoV-2 infection (Sadoff et al., 2022).

Potential observed confounders include age, participant’s region of residence, and a pre-defined risk score for acquiring symptomatic COVID-19 infection, while an individual’s immunity level is likely a strong unmeasured confounder. The ENSEMBLE trial measured

strong proxies of the immune system level that can serve as negative controls: the anti-spike binding antibody level at D29 as negative control treatment and the burden of respiratory symptoms unrelated to SARS-CoV-2 infection during trial follow-up as negative control outcome.

Measuring anti-spike binding antibodies is a more accessible alternative to assessing neutralizing antibodies, as it does not require biosafety level 2 or 3 facilities. However, the anti-spike binding antibody assay has a key limitation: it only detects antibodies that bind to the SARS-CoV-2 spike protein without distinguishing those with actual neutralizing activity. Thus, while it can identify the presence of antibodies, it cannot measure their capacity to neutralize the virus (Sholukh et al., 2020). In our analysis, we treat anti-spike binding antibodies as suitable negative control treatments. This approach aligns with an immunological model where binding to the spike protein is a necessary intermediate step in antibody-mediated protection, but only neutralizing antibodies directly block host cell infection. Under this model, anti-spike binding antibodies should not have an independent effect on the outcome once neutralizing antibodies are accounted for.

Similarly, SARS-CoV-2-specific antibodies should not protect against other respiratory infections, making non-COVID respiratory symptoms suitable negative control outcomes. The burden of these symptoms serves as a proxy for an individual’s general susceptibility to respiratory infections, which may correlate with their underlying immunity level without being directly affected by SARS-CoV-2 antibodies.

3.3 Target of Inference and Identification

Consider an observational study comprising n individuals randomly selected from a target population. For each subject $i = 1, \dots, n$, we observe a continuous exposure A_i , an outcome Y_i , and m potential confounders $L_i = (L_{i,1}, \dots, L_{i,m})$. Let $Y_i(a)$ be the counterfactual

outcome when subject i receives exposure level $A_i = a$. We are interested in estimating

$$\psi_0 = \mathbb{E} [Y_i \{q(A_i, L_i)\}],$$

where $q : \text{supp}(A, L) \rightarrow \mathbb{R}$ is a given function. For instance, if $q(a, l) = a + \delta$, then ψ_0 is the expected value of the outcome in a hypothetical scenario where every subject receives an exposure that is δ units higher than the exposure they actually received.

In our motivating application, the sample consists of the participants randomized to the vaccine arm in the ENSEMBLE trial who actually received the vaccine and were sero-negative at enrollment (i.e., had no detectable levels of antibodies against the SARS-CoV-2 virus). The variable A is the vaccine-induced neutralizing antibody levels in D29 on a logarithmic scale of base 10, Y is the indicator of symptomatic COVID-19 infection between 7 and 210 days after D29, and L is the vector of observed variables deemed as potential confounders as described earlier. Thus, in this application, when $q(a, l) = a + \delta$, ψ_0 corresponds to the probability of symptomatic COVID-19 in a hypothetical world where the levels of neutralizing antibodies of each participant are 10^δ times higher than those observed in the trial. The shift-intervention-curve $\delta \rightarrow \mathbb{E} [Y_i(A_i + \delta)]$ provides a useful perspective for assessing the causal effect of neutralizing antibody levels on the outcome. Unfortunately when, as in the ENSEMBLE study, the support of A is bounded, we cannot identify, and therefore hope to find consistent estimators of $\mathbb{E} [Y_i(A_i + \delta)]$ for any $\delta \neq 0$ even in the favorable scenario in which the measured covariates, say denoted with L , suffice to control for confounding between A and Y . To see this, suppose $[c, d]$ is the support of A , and write:

$$\mathbb{E}[Y(A + \delta)] = \int p(l) dl \int_c^d \mathbb{E}[Y(a + \delta) \mid A = a, L = l] p(a \mid l) da.$$

When L suffices to control for confounding, we can infer the counterfactual mean $\mathbb{E}[Y(A +$

$\delta) \mid A = a, L = l]$ from the observed data for any a and $a + \delta$ in $[c, d]$ since

$$\begin{aligned} \mathbb{E}[Y(A + \delta) \mid A = a, L = l] &= \mathbb{E}[Y(a + \delta) \mid A = a, L = l] \\ &= \mathbb{E}[Y(a + \delta) \mid A = a + \delta, L = l] \\ &= \mathbb{E}[Y \mid A = a + \delta, L = l]. \end{aligned}$$

where the third equality follow if, as we assume throughout, $Y(A) = Y$, and the second follows by $Y(a^*)$ conditionally independent of A given L . However, if $a \in [c, d]$ but $a + \delta \notin [c, d]$, the conditional mean $\mathbb{E}[Y \mid A = a + \delta, L = l]$ does not exist and the conditional counterfactual mean cannot be inferred from the observed data.

The preceding discussion illustrates the important fact that for arbitrary q , $\mathbb{E}[Y\{q(A, L)\}]$ is not identified by the observed data if there exists a in the support of A such that $q(a, l)$ falls outside the support. To address this complication, two strategies can be employed. The first is to restrict the population over which we compute the mean of $Y\{q(A, L)\}$. The second, is to change the policy in a meaningful way so that for the new policy, say q^* , $q^*(a, l)$ always lies within the support of A . The following examples illustrate these strategies for policies of the form $q(a, l) = a + \delta$.

Example 3.1 (First strategy: Restrict the population). *For a range of values of δ of interest, compute the mean of $Y\{q(A, L)\}$ over a subpopulation defined by $(A, L) \in \mathcal{S}$, where $\mathcal{S}$ is a subset of the support of (A, L) satisfying the condition that $(a + \delta, l)$ is in the support of (A, L) whenever $(a, l) \in \mathcal{S}$. For example, if the conditional support of A given L is constant and equal to $[c, d]$, and $\delta \in [0, 10]$, we can choose $\mathcal{S} = [c, d - 10] \times \text{supp}(L)$, provided $c < d - 10$.*

Example 3.2 (Second strategy: Change the target policy). *Suppose that for each l in the support of L , the support of A is a known finite interval which we denote with $[c(l), d(l)]$.*

Then, we may consider target policies of the form

$$q^\varepsilon(a, l) = \begin{cases} a + \delta & \text{if } a \in [c(l), d(l) - \delta - \varepsilon(l)], \\ a + \frac{\delta}{\delta + \varepsilon(l)} \{d(l) - a\} & \text{if } a \in [d(l) - \delta - \varepsilon(l), d(l)]. \end{cases}$$

for a specified positive function $l \rightarrow \varepsilon(l)$ such that for all $l \in \text{supp}(L)$, $c(l) < d(l) - \delta - \varepsilon(l)$. The policy $q^\varepsilon(a, l)$ shifts each individual's observed exposure value by δ if the observed exposure is lower than $d(l) - \delta - \varepsilon(l)$, and by $\frac{\delta}{\delta + \varepsilon(l)} \{d(l) - a\}$ otherwise. Notably, for individuals with an observed exposure value within $\delta + \varepsilon(l)$ units from $d(l)$, the shift decreases in proportion to the exposure value. For individuals with an observed exposure value of $d(l)$, the policy makes no shift.

The tradeoffs between these two strategies are evident. The first strategy focuses on simple and interpretable shift policies but restricts the analysis to a truncated subpopulation, limiting generalizability. In contrast, the second strategy evaluates effects across the entire population but introduces slightly more complex and less intuitive policies.

To minimize distracting notational complexity, in the main text we shall focus on target parameters defined using the second strategy, alternatively, parameters as in the first strategy in which $\mathcal{S}$ is the full support of (A, L) . The more general case corresponding to $\mathcal{S}$ being a specified strict subset of the support of (A, L) is deferred to Appendix B.1.

3.3.1 Proximal Identifying Assumptions

In studies of correlates of protection, as in observational studies more broadly, the estimation of counterfactual MTP means is often hindered by unmeasured confounders, such as the unobserved immune status of individuals in the aforementioned studies.

Building on the foundational contributions of Miao et al. (2018); Tchetgen et al. (2020); and Cui et al. (2024) (see also Kallus et al. (2021)), we demonstrate in this section that,

under certain regularity conditions, the counterfactual MTP mean ψ_0 becomes identifiable if, in addition to baseline covariates L , we also measure a negative control treatment Z and a negative control outcome W . Heuristically, a negative control treatment Z is a variable influenced by the unmeasured confounder U that, either has no direct effect on the outcome Y or has an effect on Y that is entirely mediated through the exposure A . Similarly, a negative control outcome W is a variable influenced by the unmeasured confounder U , and potentially by Y , but which is unaffected by the treatment X and the negative control treatment Z . These relationships are depicted in Figure 3.1.

We assume that $(U_i, L_i, Z_i, W_i, A_i, Y_i^c)$ where $Y_i^c := \{Y_i(a) : a \in \text{support of } A\}$ are n i.i.d. copies of a random vector (U, L, Z, W, A, Y^c) , and for each subject i , we observe $O_i = (L_i, Z_i, W_i, A_i, Y_i)$ where $Y_i = Y_i(A_i)$ as stated in Assumption 3.1. That is, we never observe U and, of the vector of potential outcomes under all possible treatment doses, we only observe the outcome corresponding to the dose actually received. We make the following assumptions:

Assumption 3.1 (Consistency). $Y(A) = Y$.

Assumption 3.2 (Randomization of policy related populations). *For all $(a, l, u) \in \text{supp}(A, L, U)$ and $a' = q(a, l)$, it holds that $\text{law}\{Y(a')|A = a, L = l, U = u\} = \text{law}\{Y(a')|A = a', L = l, U = u\}$.*

Assumption 3.3 (Common support). *For all $(l, u) \in \text{supp}(L, U)$, $\text{supp}(A|L = l, U = u) \equiv \text{supp} A|L = l$.*

Assumption 3.4 (Negative control treatment). $Z \perp Y|A, L, U$.

Assumption 3.5 (Negative control outcome). $(A, Z) \perp W|U, L$.

Above and hereafter, for any random vectors V and L , $\text{supp}(V)$ and $\text{supp}(V|L = l)$ denote the supports of the distributions of V and of V given $L = l$ respectively. Also $W \perp R|S$ stands for W conditionally independent of R given S .

The first assumption is standard in causal inference. The second assumption is tantamount to the exchangeability between the population with the given observed exposure and the population whose observed exposure matches the exposure the first population would have under the policy. This condition must hold for every possible observed exposure level. Assumption 3 is analogous to the positivity assumption of the propensity score in the proximal causal inference literature for binary treatments, which similarly requires positivity conditional on unmeasured confounders (Miao et al., 2018; Tchetgen et al., 2020; Kallus et al., 2021; Ghassami et al., 2022; Cui et al., 2024). This assumption cannot be empirically tested because U is unobserved. Assumptions 3.4 and 3.5 are implied by the negative control relationships depicted in Figure 3.1. In our motivating example, anti-spike binding antibodies act as negative control treatments within an immunological framework where they influence the outcome solely through their neutralizing activity. Meanwhile, the burden of non-COVID-19 respiratory symptoms serves as a negative control outcome, as these symptoms should not be affected by SARS-CoV-2 antibodies.

We also impose the following assumption on the policy:

Assumption 3.6 (Monotone and smooth policy). *For each $l \in \text{supp}(L)$, $q(\cdot, l)$ is strictly monotone and differentiable almost everywhere with respect to the distribution of A given $L = l$.*

This assumption ensures that for each l , the policy map $q : a \mapsto q(a, l)$ has an inverse that is almost everywhere differentiable with respect to the distribution of A given $L = l$. Both the shift intervention policy and its modified version, as discussed in Example 3.2, satisfy this assumption.

Lastly, following Miao et al. (2018) and Cui et al. (2024), we make assumptions that allow us to leverage the negative controls to address the unobserved confounding. Specifically, as established in Theorem 3.2 below, identifiability of the counterfactual MTP mean is possible

if, in addition to Assumptions 3.1–3.6, one of the Assumptions 3.7 or 3.8 below holds.

Hereafter, for any random vector V , $\mathcal{L}^2(V)$ denotes the space of functions of V with finite second moment. Additionally, hereafter we assume that for every l in the support of L , the support of the distribution of A given $L = l$ is an interval which we denote as $[c(l), d(l)]$.

Assumption 3.7 (Existence of latent outcome bridge and observed treatment bridge functions).

- i) There exists $h_0 \in \mathcal{L}^2(A, L, W)$, hereafter called “latent outcome bridge function”, solving for all $(a, l, u) \in \text{supp}(A, L, U)$, the integral equation

$$\mathbb{E}(Y|A = a, L = l, U = u) = \int h(a, l, w) p_{W|A, L, U}(w|a, l, u) dw. \quad (3.1)$$

- ii) There exists $g_0 \in \mathcal{L}^2(A, L, Z)$, hereafter called “observed treatment bridge function”, such that for all $(a, l, w) \in \text{supp}(A, L, W)$ satisfying that the map $a' \mapsto q^{-1}(a', l)$ is differentiable at $a' = a$ and $a \in q([c(l), d(l)], l)$, g_0 solves the integral equation

$$\alpha_0(a, l, w) = \int g(a, l, z) p_{Z|A, L, W}(z|a, l, w) dz, \quad (3.2)$$

where

$$\alpha_0(a, l, w) := I\{a \in q([c(l), d(l)], l)\} \frac{dq^{-1}(a, l)}{da} \frac{p_{A|L, W}\{q^{-1}(a, l)|l, w\}}{p_{A|L, W}(a|l, w)}.$$

and $g_0(a, l, z) = 0$ otherwise.

Assumption 3.8 (Existence of latent treatment bridge and observed outcome bridge functions).

- i) There exists $g_0 \in \mathcal{L}^2(A, L, Z)$, hereafter called “latent treatment bridge function”, such

that for all $(a, l, u) \in \text{supp}(A, L, U)$ satisfying that the map $a' \mapsto q^{-1}(a', l)$ is differentiable at $a' = a$ and $a \in q([c(l), d(l)], l)$, g_0 solves the integral equation

$$\alpha_0(a, l, u) = \int g(a, l, z) p_{Z|A, L, U}(z|a, l, u) dz, \quad (3.3)$$

where

$$\alpha_0(a, l, u) := I\{a \in q([c(l), d(l)], l)\} \frac{dq^{-1}(a, l)}{da} \frac{p_{A|L, U}\{q^{-1}(a, l)|l, u\}}{p_{A|L, U}(a|l, u)} \quad (3.4)$$

and $g_0(a, l, z) = 0$ otherwise.

ii) There exists $h_0 \in \mathcal{L}^2(A, L, W)$, hereafter called “observed outcome bridge function”, solving for all $(a, l, z) \in \text{supp}(A, L, Z)$, the integral equation

$$\mathbb{E}(Y|A = a, L = l, Z = z) = \int h(a, l, w) p_{W|A, L, Z}(w|a, l, z) dw. \quad (3.5)$$

The existence of the bridge functions postulated in Assumptions 3.7 or 3.8, encodes the precise sense in which the negative control variables suffice to capture the variability of the unmeasured U in order to control for unmeasured confounding. To gain some intuition, consider Assumption 3.7 (a similar interpretation applies to Assumption 3.8 but with the roles of W and Z reversed). Following Miao et al. (2018) and Cui et al. (2024), and using Piccard’s theorem (see Theorem 15.18 in Kress (2014)), we show in Appendix B.1.2 that equation (3.1) has a solution if for all (a, l) :

$$\eta(a, l, U = u) = 0 \quad \forall u \quad \text{whenever} \quad \mathbb{E}\{\eta(A, L, U)|A = a, L = l, W = w\} = 0 \quad \forall w. \quad (3.6)$$

This condition is sufficient when W and U are discrete and finitely valued; otherwise, it must

be combined with a regularity condition provided in Appendix B.1.2. Condition (3.6) is often referred to as a completeness property. It asserts that, within any fixed level of (A, L) , no function of U can be orthogonal to the information provided by W . In other words, W must be sufficiently informative about U such that any transformation of U that averages to a constant when conditioned on W must itself be constant. In particular, this implies that, within any fixed level of (A, L) , the first, second, and all higher moments of U vary with W . Put differently, W has predictive value for every possible direction of variation in U . Similarly, it can be shown that under regularity conditions, equation (3.2) has a solution if for all (a, l)

$$\eta(a, l, W = w) = 0 \quad \forall w \quad \text{whenever} \quad \mathbb{E}\{\eta(A, L, W) | A = a, L = l, Z = z\} = 0 \quad \forall z$$

which indicates that, conditional on (A, L) , Z must have predictive value for every possible direction of variation in W . Under the conditional independence postulated by Assumption 3.5, this is only possible due to dependence of U and Z . For additional insights, see Miao et al. (2018); Tchetgen et al. (2020); Shi et al. (2020); Kallus et al. (2021), and Cui et al. (2024).

The following result establishes that for negative control variables, latent bridge functions are also observed bridge functions.

Theorem 3.1.

- i) [Kallus et al. (2021)] Under Assumptions 3.4 and 3.5, any latent outcome bridge function h_0 , i.e. satisfying the conditions in Assumption 3.7 (i), is also an observed outcome bridge function, i.e. it satisfies the conditions in Assumption 3.8 (ii).
- ii) Under Assumption 3.5, any latent treatment bridge function g_0 is also an observed treatment bridge function.

With these premises, building on the approaches of Miao et al. (2018); Tchetgen et al. (2020); Kallus et al. (2021); Ghassami et al. (2022), and Cui et al. (2024), we are now in position to derive the identification result for the MTP counterfactual mean. Hereafter, we define for any $h \in \mathcal{L}^2(A, L, W)$ and $g \in \mathcal{L}^2(A, L, Z)$,

$$\phi(O; h, g) := h\{q(A, L), L, W\} + g(A, L, Z)\{Y - h(A, L, W)\}.$$

Theorem 3.2. *Suppose Assumptions 3.1–3.6 hold. If, in addition, either Assumption 3.7 holds or Assumption 3.8 holds, then ψ_0 is identified by the observed data and it satisfies all of the following:*

$$\psi_0 = \mathbb{E}[h_0\{q(A, L), L, W\}] \tag{3.7}$$

$$= \mathbb{E}\{Y g_0(A, L, Z)\} \tag{3.8}$$

$$= \mathbb{E}\{\phi(O; h^\dagger, g^\dagger)\} \tag{3.9}$$

where h_0 is an observed outcome bridge function, g_0 is an observed treatment bridge function and, either $h^\dagger$ is an observed outcome bridge function or $g^\dagger$ is an observed treatment bridge function but not necessarily both functions are observed bridge functions.

The identifying expressions in Theorem 3.2 parallel those that hold when U is observed. In that case, ψ_0 equals any of the three expressions in Theorem 3.2, with $h_0(a, l, w)$ replaced by $\mathbb{E}(Y|A = a, L = l, U = u)$, $g_0(a, l, z)$ by $\alpha_0(a, l, u)$ as defined in (3.4), and $h^\dagger$ and $g^\dagger$ by functions of $e(a, l, u)$ and $\alpha(a, l, u)$ such that $\alpha = \alpha_0$ or $e(a, l, u) = \mathbb{E}(Y|A = a, L = l, U = u)$ (Díaz Muñoz and Van Der Laan, 2012; Haneuse and Rotnitzky, 2013). In fact, when $Z = W$, the three representations of the counterfactual MTP mean in Theorem 3.2 are exactly equal to the corresponding representations in the absence of unmeasured confounding. We refer to

the expressions in the right hand sides of (3.7), (3.8) and (3.9) as outcome regression (OR), density-quotient-weighted (DQW) and doubly-robust (DR), respectively.

The results in Theorem 3.2 also align with those in the proximal causal inference literature for other target parameters. For example, Miao et al. (2018), Kallus et al. (2021) and Cui et al. (2024) established the identifiability of the counterfactual mean in a hypothetical scenario where everyone receives the same treatment, under conditions similar to the assumptions in Theorem 3.2, but with α_0 replaced by the inverse of the propensity score. Notably, as in their work, the identifying expressions in Theorem 3.2 remain valid even if equations (3.2) and/or (3.5) have multiple solutions, provided g_0 and h_0 are any solutions to these equations, respectively.

3.4 Estimation

Theorem 3.2 suggests three potential estimators for the target parameter ψ_0 , namely

$$\begin{aligned}\hat{\psi}^{OR} &= \mathbb{E}_n \left[\hat{h} \{q(A, L), L, W\} \right], \\ \hat{\psi}^{DQW} &= \mathbb{E}_n [Y \hat{g}(A, L, Z)],\end{aligned}$$

and

$$\hat{\psi}^{DR} = \mathbb{E}_n \left[\phi(O; \hat{h}, \hat{g}) \right],$$

where $\hat{g}$ and $\hat{h}$ are estimators of some solutions g_0 and h_0 of equations (3.2) and (3.5) respectively, and throughout $\mathbb{E}_n[\cdot]$ denotes the empirical mean operator, i.e. $\mathbb{E}_n[t(O)] = \frac{1}{n} \sum_{i=1}^n t(O_i)$ for n i.i.d. copies $O_1, \dots, O_n$ of O .

The estimators $\hat{\psi}^{OR}$ and $\hat{\psi}^{DQW}$ generally fail to achieve $\sqrt{n}$ -consistency unless $\hat{h}$ and $\hat{g}$ converge to some respective observed bridge functions at parametric rates. However, this requires correctly specified parametric models for the bridge functions—an often unrealis-

tic assumption. In contrast, as established in Theorem 4.1 below, the estimator $\hat{\psi}^{DR}$ can achieve $\sqrt{n}$ -consistency and asymptotic normality, under regularity conditions, even when the bridge functions are estimated non-parametrically. For these reasons, hereafter we focus exclusively on $\hat{\psi}^{DR}$ with bridge functions estimated non-parametrically. In particular, we adapt recently proposed non-parametric methods to estimate the bridge functions that are specific to our problem. Nevertheless, for completeness, in the Appendix B.3.3, we provide unbiased estimating equations for parametric bridge function estimation. Furthermore, we describe the construction of doubly-robust confidence intervals, i.e. intervals whose asymptotic coverage probability equals the nominal one as long as one parametric bridge estimator converges in probability to a solution of the corresponding observed bridge equation, regardless of which one. Following the modern theory of debiased machine learning (Robins et al., 2008; Chernozhukov et al., 2018), we describe a cross-fitted version of the estimator $\hat{\psi}^{DR}$. For both $\hat{\psi}^{DR}$ and its cross-fitted version to converge to a normal distribution with mean zero at a $\sqrt{n}$ -rate, a certain random quantity—defined later—must be of order $o_p(n^{-1/2})$. The estimator $\hat{\psi}^{DR}$ further requires that the estimators $\hat{h}$ and $\hat{g}$, along with the observed bridge functions h_0 and g_0 they are converging to, lie within function classes subject to specific size constraints, such as Donsker classes. In contrast, this condition is not necessary for the cross-fitted version of $\hat{\psi}^{DR}$ to achieve $\sqrt{n}$ convergence. We focus on the cross-fitted estimator because this relaxation broadens the set of data-generating processes under which the estimator of the target parameter is asymptotically normal.

To construct the cross-fitted DR estimator, we partition the sample into K approximately equal-sized subsets, denoted $I_1, \dots, I_K$, where K is a small number, such as $K = 5$. For each $k \in \{1, \dots, K\}$, using data from all subsamples except I_k , we obtain estimators $\hat{h}^{(-k)}$ and

$\hat{g}^{(-k)}$ for the bridge functions h_0 and g_0 . The cross-fitted DR estimator is then calculated as

$$\hat{\psi}_{CF}^{DR} = \frac{1}{n} \sum_{k=1}^K \sum_{i: O_i \in I_k} \phi \{O_i; \hat{h}^{(-k)}, \hat{g}^{(-k)}\}.$$

The next theorem establishes results on the asymptotic behavior of $\hat{\psi}_{CF}^{DR}$. Hereafter, for any function f of the random vector O , we let $\|f\|_\infty := \sup_{o \in \text{supp}(O)} |f(o)|$ and $\|f\|_2 := [\mathbb{E}\{f(O)^2\}]^{1/2}$ respectively. Note that $\|\cdot\|_2$ is the 2-norm in the space $\mathcal{L}^2(O)$.

Theorem 3.3. *Suppose that the Assumptions of Theorem 3.2 hold. Let h_0 and g_0 denote observed outcome and treatment bridge functions respectively. If for some $B > 0$, $|\alpha_0(A, L, W)| \leq B$, $|Y| \leq B$ and for each k , $\|\hat{h}^{(-k)} - h_0\|_2 = o_p(1)$, $\|\hat{g}^{(-k)} - g_0\|_2 = o_p(1)$ and either $\|h_0\|_\infty + \|\hat{g}^{(-k)}\|_\infty \leq B$ or $\|\hat{h}^{(-k)}\|_\infty + \|g_0\|_\infty \leq B$, then*

$$\sqrt{n} \{ \hat{\psi}_{CF}^{DR} - \psi_0 \} = \sqrt{n} \mathbb{E}_n \{ \phi(O; h_0, g_0) - \psi_0 \} + \sqrt{n} R_n + o_p(1).$$

where $R_n = \sum_{k=1}^K \frac{n_k}{n} \int \{ \hat{h}^{(-k)}(a, l, w) - h_0(a, l, w) \} \{ g_0(a, l, z) - \hat{g}^{(-k)}(a, l, z) \} dP(a, l, w, z)$ and n_k denoting the cardinality of I_k . In particular, if $\mathbb{E} \{ \phi(O; h_0, g_0)^2 \} < \infty$ and $R_n = o_p(n^{-1/2})$, then

$$\sqrt{n} \{ \hat{\psi}_{CF}^{DR} - \psi_0 \} \xrightarrow{d} \mathcal{N}(0, \tau^2), \quad (3.10)$$

where $\tau^2 = \mathbb{E} [\{ \phi(O; h_0, g_0) - \psi_0 \}^2]$.

In the next section we provide non-parametric estimators $\hat{h}^{(-k)}$ and $\hat{g}^{(-k)}$ that, under certain regularity conditions, are $\mathcal{L}^2$ -norm for the observed bridge functions that minimize the norm induced by the function class used for estimation, among all observed bridge functions lying in that class. These estimators yield $R_n = o_p(n^{-1/2})$, thus resulting in $\hat{\psi}_{CF}^{DR}$

satisfying (3.10). The expression for τ^2 suggests estimating it with

$$\hat{\tau}^2 = \frac{1}{n} \sum_{k=1}^K \sum_{i: O_i \in I_k} \left[\phi \left\{ O_i; \hat{h}^{(-k)}, \hat{g}^{(-k)} \right\} - \hat{\psi}_{CF}^{DR} \right]^2$$

Under regularity conditions, $\hat{\tau}^2$ is a consistent estimator of τ^2 and consequently, the Wald interval $\hat{\psi}_{CF}^{DR} \pm n^{-1/2} \Phi^{-1}(1 - \xi/2) \hat{\tau}$, where $\Phi^{-1}(1 - \xi/2)$ is the $1 - \xi/2$ quantile of the standard normal distribution, covers ψ_0 with probability that converges to $1 - \xi$.

3.4.1 Estimation of the Bridge Functions

To estimate the observed bridge functions, we employ kernel regularized adversarial stabilized estimators, a class of adversarial methods recently introduced by [Dikkala et al. \(2020\)](#) and [Ghassami et al. \(2022\)](#). In Chapter 4, we review these methods and develop a novel convergence analysis for these estimators, establishing finite-sample bounds on both the root-mean squared error and the weak error.

The estimation procedure is as follows. Define

$$\Psi_{H,n}^{(-k)}(h, g) := \mathbb{E}_n^{(-k)} [g(A, L, Z) \{Y - h(A, L, W)\} - g(A, L, Z)^2],$$

$$\Psi_{G,n}^{(-k)}(g, h) := \mathbb{E}_n^{(-k)} [h \{q(A, L), L, W\} - I_q(A, L)h(A, L, W)g(A, L, Z) - I_q(A, L)h(A, L, W)^2],$$

where $I_q(a, l) := I\{a \in q([c(l), d(l)], l)\}$ and $\mathbb{E}_n^{(-k)}$ denotes the empirical mean over all the observations excluding those in fold I_k . Let $\mathcal{H}, \mathcal{H}' \subseteq \mathcal{L}^2(A, L, W)$ and $\mathcal{G}, \mathcal{G}' \subseteq \mathcal{L}^2(A, L, Z)$ denote function classes, each of them included, in possibly distinct, reproducing kernel Hilbert spaces (RKHSs) specified by the data analyst. Let $\|\cdot\|_{\mathcal{H}}, \|\cdot\|_{\mathcal{H}'}, \|\cdot\|_{\mathcal{G}}$, and $\|\cdot\|_{\mathcal{G}'}$ denote the norms of the RKHSs where $\mathcal{H}, \mathcal{H}', \mathcal{G}$, and $\mathcal{G}'$ are included. For each $k \in \{1, \dots, K\}$, we compute the estimators $\hat{h}^{(-k)}$ and $\hat{g}^{(-k)}$ of the bridge functions, using all the observations

but those in fold I_k , as follows:

$$\hat{h}^{(-k)} = \arg \min_{h \in \mathcal{H}} \max_{g \in \mathcal{G}'} \left\{ \Psi_{H,n}^{(-k)}(h, g) - \lambda_{\mathcal{G}'} \|g\|_{\mathcal{G}'}^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2 \right\}, \quad (3.11)$$

$$\hat{g}^{(-k)} = I_q(A, L) \cdot \arg \min_{g \in \mathcal{G}} \max_{h \in \mathcal{H}'} \left\{ \Psi_{G,n}^{(-k)}(g, h) - \lambda_{\mathcal{H}'} \|h\|_{\mathcal{H}'}^2 + \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2 \right\}, \quad (3.12)$$

where $\lambda_{\mathcal{H}}$, $\lambda_{\mathcal{H}'}$, $\lambda_{\mathcal{G}}$, and $\lambda_{\mathcal{G}'}$ are positive tuning parameters. By restricting attention to function classes within RKHSs, the estimators $\hat{h}^{(-k)}$ and $\hat{g}^{(-k)}$ admit closed-form expressions. Before presenting these expressions, a few remarks are in order.

Remark 3.1. *The min-max optimization problems (3.11) and (3.12) defining the estimators include penalization terms in both the outer and inner optimization components. The outer terms, $\lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2$ and $\lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2$, known as Tikhonov regularization terms, primarily serve to control the complexity of the solutions and enhance their stability. Tikhonov regularization is a standard technique for solving Type I integral equations like (3.2) and (3.5), which are ill-posed because their solutions are highly sensitive to small perturbations in the equations' left-hand sides. The inner terms, $\lambda_{\mathcal{G}'} \|g\|_{\mathcal{G}'}^2$ and $\lambda_{\mathcal{H}'} \|h\|_{\mathcal{H}'}^2$, are included solely for computational tractability. Penalizing with RKHS norms—rather than, for example, empirical $\mathcal{L}^2(\mathbb{P}_n)$ norms—avoids the inversion of nearly ill-posed matrices and ensures numerically stable computations. See Chapter 4 for further discussion.*

Remark 3.2. *The estimators $\hat{h}^{(-k)}$ and $\hat{g}^{(-k)}$ converge to the observed bridge functions with minimum RKHS norm, denoted by h_0 and g_0 , within their respective function classes. The theoretical bounds on their rates of convergence rely on a “source condition” imposed on these minimum-norm solutions, which relates their smoothness to the smoothing properties of the associated conditional expectation operators. In particular, under regularity conditions and with $\lambda_{\mathcal{H}}$ and $\lambda_{\mathcal{G}}$ selected to converge to zero at appropriate rates, the upper error bounds established in Chapter 4 imply that the remainder term R_n is $o_p(n^{-1/2})$, as required for the*

asymptotic normality of the MTP mean estimator. Importantly, this theoretical analysis remains valid regardless of whether multiple solutions to the observed bridge equations exist, as the estimators consistently target the minimum RKHS norm solutions. A more detailed treatment of the source conditions required on h_0 and g_0 is provided in Chapter 4.

Remark 3.3. The selection of the tuning parameters in practice remains an open problem, the investigation of which is beyond the scope of this work. While a natural approach is to select them by cross-validation, a key challenge lies in the choice of the risk estimator to minimize. See Section 3.5 for the strategy we used in the simulations and data analysis to select the tuning parameters.

We now derive the close form expressions for $\hat{h}^{(-k)}$ and $\hat{g}^{(-k)}$. To simplify notation, we shall omit the superscript $(-k)$ from the estimators of the bridge functions and denote with n the size of the sample from which they are computed. Let $K_{\mathcal{H}}$, $K_{\mathcal{H}'}$, $K_{\mathcal{G}}$, and $K_{\mathcal{G}'}$ denote the kernel functions associated with the RKHSs and let $K_{\mathcal{H},n}$, $K_{\mathcal{H}',n}$, $K_{\mathcal{G},n}$, and $K_{\mathcal{G}',n}$ be the $n \times n$ matrices whose (i, j) -th entry corresponds to the value of the respective kernel function evaluated at observations i and j . By the “Representer Theorem” (Schölkopf et al., 2001), $\hat{h}$ and $\hat{g}$ admit closed form expressions which are of the form:

$$\begin{aligned}\hat{h}(a, l, w) &= \sum_{i=1}^n \gamma_i K_{\mathcal{H}} \{(a, l, w), (A_i, L_i, W_i)\}, \\ \hat{g}(a, l, z) &= I\{x \in q([c(l), d(l)], l)\} \cdot \sum_{i=1}^n \theta_i K_{\mathcal{G}} \{(a, l, z), (A_i, L_i, Z_i)\},\end{aligned}$$

for some constants γ_i and θ_i , $i = 1, \dots, n$. In Appendix B.2.1 we show that $\gamma := (\gamma_1, \dots, \gamma_n)$ and $\theta := (\theta_1, \dots, \theta_n)$ are given by:

$$\gamma = (K_{\mathcal{H},n} \Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} K_{\mathcal{H},n} + n^2 \lambda_{\mathcal{H}} K_{\mathcal{H},n})^{-1} K_{\mathcal{H},n} \Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} \widetilde{Y}_n,$$

$$\theta = (K_{\mathcal{G},n}J_n\Gamma_{\mathcal{H}'}K_{\mathcal{H}',n}J_nK_{\mathcal{G},n} + n^2\lambda_G K_{\mathcal{G},n})^{-1} K_{\mathcal{G},n}J_n\Gamma_{\mathcal{H}'}\widetilde{K}_{\mathcal{H},n}^T 1_n.$$

Here, $\widetilde{Y}_n = (Y_1, \dots, Y_n)^T$ is the vector of observed outcomes, $1_n = (1, \dots, 1)^T$ is the vector of ones, J_n is the $n \times n$ diagonal matrix with diagonal values $I\{A_i \in q([c(L_i), d(L_i)], L_i)\}$ for $i = 1, \dots, n$, $\widetilde{K}_{\mathcal{H},n}$ is the $n \times n$ matrix with (i, j) -entry equal to $K_{\mathcal{H}'}\{(q(A_i, L_i), L_i, W_i), (A_j, L_j, W_j)\}$. $\Gamma_{\mathcal{G}'} = 4^{-1}[n^{-1}K_{\mathcal{G}',n} + \lambda_{\mathcal{G}'}I_n]^{-1}$ and $\Gamma_{\mathcal{H}'} = 4^{-1}[n^{-1}K_{\mathcal{H}',n}J_n + \lambda_{\mathcal{H}'}I_n]^{-1}$ where I_n is the $n \times n$ identity matrix.

3.5 Numerical Experiments

We conducted several simulation experiments to assess the finite-sample performance of $\hat{\psi}_{CF}^{DR}$ under varying levels of correlation between the negative controls (Z and W) and the unmeasured confounder (U), given observed confounders (L). For comparison, we evaluated two existing estimators of the MTP mean that assume no unmeasured confounding: the augmented inverse probability weighted estimator (Non-proximal AIPW) and the targeted minimum loss-based estimator (Non-proximal TMLE) (Díaz Muñoz and Van Der Laan, 2012; Díaz et al., 2023). Each simulation experiment had 500 replications. We repeated the experiments for sample sizes $n = 750, 1500, 3000$, and 6000.

Let $\mathcal{TN}_{[b_1, b_2]}(\mu, \sigma^2)$ denote the $\mathcal{N}(\mu, \sigma^2)$ distribution truncated within the interval $[b_1, b_2]$. Define $\mathcal{TN}_c(\mu, \sigma^2)$ as its centered version truncated at three standard deviations around the mean, i.e., $\mathcal{TN}_{[\mu-3\sigma, \mu+3\sigma]}(\mu, \sigma^2)$. For each $i = 1, \dots, n$, we generated $L_i \sim \mathcal{TN}_c(0, 1)$, $U_i \sim \mathcal{TN}_c(0.5L_i, 1)$, $Z_i \sim \mathcal{TN}_c(0.2L_i + \beta_z U_i, 1)$, $W_i \sim \mathcal{TN}_c(0.5L_i + \beta_w U_i, 1)$, $A_i \sim \mathcal{TN}_{[-2, 2]}(0.3L_i + U_i, 1)$, and $Y_i \sim \text{Bernoulli}(p_i)$ where $p_i = \{1 + \exp(1 - 0.5L_i + U_i + 1.5A_i + 0.75A_i^2)\}^{-1}$. We generated data across nine scenarios by combining $\beta_z \in \{2, 1, 0.5\}$ (corresponding to $\text{Cor}(Z, U|L) = 0.894, 0.707$, and 0.447) and $\beta_w \in \{-2, -1, -0.5\}$ (which corresponds to $\text{Cor}(W, U|L) = -0.894, -0.707$, and -0.447). The target parameter was $\mathbb{E}[Y\{q(A)\}]$ with

$q(a) = a + 0.4I\{-2 \leq a \leq 0.6\} + \frac{0.4(2-a)}{1.4}I\{0.6 < a \leq 2\}$, the shift intervention policy in Example 3.2, with $\delta = 0.4$ and $\varepsilon = 1$. Assuming (U, L) suffice to control for confounding, then $\mathbb{E}[Y\{q(A)\}] = 0.2512$ under any of the data generating scenarios we considered. Additionally, the assumptions of Theorem 3.2 hold. The expressions for outcome and treatment bridge functions are provided in Appendix B.3.1.

We evaluated $\hat{\psi}_{CF}^{DR}$ computed with bridge functions estimated as in (3.11) and (3.12), using Gaussian RKHSs. We standardized A , L , Z , and W before implementing the estimation of the bridge functions in order to prevent any of these variables from disproportionately influencing the estimation due to differences in scale. We used tuning parameters $\lambda_{\mathcal{H}} = c_1 \left\{ \frac{\log(n')}{n'} \right\}^{1/2}$, $\lambda_{\mathcal{G}'} = c_2 \frac{\log(n')}{n'}$, $\lambda_{\mathcal{G}} = c_3 \left\{ \frac{\log(n')}{n'} \right\}^{1/2}$, and $\lambda_{\mathcal{H}'} = c_4 \frac{\log(n')}{n'}$, where $n' = n/3$, $c_1, c_3 \in \{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$ and $c_2, c_4 \in \{0.1, 1, 10, 100\}$, and fixed specified kernel bandwidths $(\sigma_{\mathcal{H}}^2, \sigma_{\mathcal{G}'}^2, \sigma_{\mathcal{G}}^2, \sigma_{\mathcal{H}'}^2)$. The bandwidth $\sigma_{\mathcal{G}'}^2$ ($\sigma_{\mathcal{H}'}^2$) was fixed at 1/4 times the median of the pairwise euclidean distance between the observed vectors (A_i, L_i, Z_i) $[(A_i, L_i, W_i)]$ (Garreau et al., 2017). The bandwidths $\sigma_{\mathcal{H}}^2$ and $\sigma_{\mathcal{G}}^2$ were also computed with the median of the pairwise Euclidean distances between observations, but multiplied by $c_5 \in \{1/4, 1/2, 1, 2, 4\}$. We used 3-fold cross-fitting, i.e. we partitioned the sample into three folds I_1, I_2 and I_3 and proceeded as explained in Section 3.4. To compute $\hat{h}^{(-1)}$, we first used fold 2 data to estimate the minimum RKHS norm solution h_0 using the various combinations of hyperparameters and bandwidths. In fold 3, we selected the configuration that minimized a specific empirical loss function detailed in Appendix B.3.2. Finally, $\hat{h}^{(-1)}$ was set as the estimator from fold 2 corresponding to the selected configuration. This procedure was similarly applied to compute $\hat{h}^{(-2)}$, $\hat{h}^{(-3)}$, and $\hat{g}^{(-k)}$ for $k = 1, 2, 3$.

For implementing the non-proximal MTP estimators of Díaz Muñoz and Van Der Laan (2012) and Díaz et al. (2023), we used the open-source R package `lmtip` with several libraries for estimation of the outcome regression and density ratio functions. In addition

to adjusting for the observed confounder L , the non-proximal estimators also adjusted for the negative controls Z and W , which correspond to precision variables under the specified data-generating mechanisms. We included these adjustments for a fair comparison with the proximal estimators.

The performance of our estimator (Proximal DR) and its competitors is summarized in Figures 3.2–3.4. In the majority of the scenarios, the estimator $\hat{\psi}_{CF}^{DR}$ exhibits low bias, and the average estimated asymptotic variance closely aligns with the empirical variance of $\hat{\psi}_{CF}^{DR}$. The coverage of the Wald confidence intervals centered at $\hat{\psi}_{CF}^{DR}$ is close to the nominal level of 0.95. Thus, the proposed estimator performed as predicted by theory across various scenarios with differing correlations between the negative controls and the unmeasured confounder, with the performance deteriorating as these correlations decrease. As expected, it outperformed conventional non-proximal methods that cannot account for unmeasured confounding.

In Appendix B.3.3, we also present results for a doubly-robust estimator that employs parametric estimators of the bridge functions. These simulations illustrate that estimation of the counterfactual MTP mean also deteriorates as the conditional correlation between the negative controls and the unmeasured confounder weakens, even when the estimation is based on correctly specified parametric models for the bridge functions.




Figure 3.2: Boxplots of the estimators from numerical experiments with data generated under varying conditional correlations between the proxies Z and W and the unmeasured confounder U , with sample sizes $n = 750, 1500, 3000$, and 6000 . The dashed line represents the true parameter value. Non-proximal estimators rely on the (incorrect) assumption of no unmeasured confounding.




Figure 3.3: Ratio of empirical variance to the mean estimated asymptotic variance for each estimator, based on numerical experiments with data generated under varying conditional correlations between the proxies Z and W and the unmeasured confounder U , across sample sizes $n = 750, 1500, 3000$, and 6000 . The dashed line indicates a variance ratio of 1.




Figure 3.4: Coverage probabilities of the confidence intervals from numerical experiments with data generated under varying conditional correlations between the proxies Z and W and the unmeasured confounder U , across sample sizes $n = 750, 1500, 3000$, and 6000 . The dashed line indicates the nominal coverage level of 0.95. Non-proximal estimators rely on the (incorrect) assumption of no unmeasured confounding.

3.6 Application to the *ENSEMBLE* Trial

We applied our methodology to investigate neutralizing antibody levels 29 days (D29) post-vaccination as correlates of protection against symptomatic COVID-19 infection in the *ENSEMBLE* trial. The analysis included 1,164 vaccine recipients who were SARS-CoV-2 seronegative at baseline, had D29 antibody measurements and came from the six participating countries (out of eight) which had at least 50 subjects with antibody measurements. These subjects were selected using a two-phase case-cohort sampling design, as outlined in Gilbert et al. (2022b). To account for this sampling design, we employed a weighted version of our estimator, as detailed in Appendix B.2.2. The outcome of interest was the binary indicator of developing symptomatic, molecularly confirmed COVID-19 infection with onset between 7 and 210 days after D29. During this follow-up period, 366 symptomatic COVID-19 cases were observed among vaccine recipients with antibody measurements.

Our analysis included the baseline covariates L : age, geographic region, and a pre-specified risk score for acquiring symptomatic COVID-19. Anti-Spike binding antibody levels were used as a negative control treatment, and the burden of non-SARS-CoV-2 respiratory symptoms during follow-up served as a negative control outcome. The burden of respiratory symptoms was calculated as the average number of symptoms per day not attributed to SARS-CoV-2, normalized by the maximum average within each country to account for seasonal variations across trial locations.

We evaluated five shift intervention policies from Example 3.2 with $\delta \in \{0.2, 0.4, 0.6, 0.8, 1\}$, $\varepsilon \equiv 1$, and constant $c(\cdot)$ and $d(\cdot)$. Neutralizing antibody levels at D29 (log10 scale) ranged from 0.3889 to 4.2175. The percentages of vaccine recipients with observed antibody levels within the range of each policy were 42.5%, 38.2%, 27.3%, 18.8%, and 11.3%, respectively. We applied our doubly robust cross-fitted estimator using the same folds, function classes, and tuning parameters as in the numerical experiments in Section 3.5.

The probability of symptomatic COVID-19 infection between 7 and 210 days after D29 for the analyzed sample was estimated to be 0.0235, (95% CI 0.0204 - 0.0265). This probability was estimated to decrease to 0.0105 (95% CI 0.0070 - 0.0140) under the policy associated with $\delta = 1$. Figure 3.5 presents the results for the policies corresponding to the values of δ that we considered. Additionally, for each considered policy q we computed estimators of

$$VE(q) := 1 - \frac{\mathbb{E}[Y\{X = 1, q(A)\}]}{\mathbb{E}[Y(X = 0)]},$$

as well as $VE := VE(Id)$. Here, $Y(X = 0)$ represents the outcome under placebo, $Y\{X = 1, q(A)\}$ represents the outcome under a vaccine formulation that induces an antibody level $q(A)$ instead of the current level A , and Id is the no-shift MTP, so that VE is the usual vaccine efficacy measure. The measure $VE(q)$ quantifies the relative reduction in the probability of symptomatic COVID-19 infection between a hypothetical world where everyone receives the vaccine and develops neutralizing antibody levels according to the specified policy q and one where everyone receives placebo. We estimated $\mathbb{E}[Y(X = 0)]$ with the sample mean of the outcomes of the 17,108 placebo recipients from the same six countries as the analyzed vaccine recipients. Among these placebo recipients, there were 805 COVID-19 cases during the same follow-up period, resulting in an observed probability of symptomatic COVID-19 infection of 0.0471 (95% CI 0.0439 - 0.0503). Combining this information with that from the analyzed vaccine recipients, the observed VE in the trial was 0.502 (95% CI 0.428 - 0.575). Under a policy that increases each participant's antibody levels by one logarithmic unit, the VE improved to 0.777 (95% CI 0.701 - 0.853). Vaccine efficacy estimates for the other policies are also presented in Figure 3.5.




Figure 3.5: Results from application to data from the ENSEMBLE trial. (Panel A) Observed and counterfactual probabilities of symptomatic COVID-19 infection and (Panel B) observed and counterfactual vaccine efficacy (VE) against symptomatic COVID-19 infection for policies that shift everyone's observed antibody levels by 0 (no shift), 0.2, 0.4, 0.6, 0.8, and 1 logarithm. The bars indicate the 95% confidence interval at each shift without any correction.

3.7 Discussion

In this chapter, we established identification conditions for causal estimands involving modified treatment policies in the presence of unmeasured confounding, assuming the availability of negative control treatment and outcome variables.

Unmeasured confounding is a significant concern in the immunobridging studies motivating our methodology, as individual immunity levels are likely important but unobserved confounders.

We studied an asymptotically normal, cross-fitted estimator that relies on nonparametric estimation of observed bridge functions. The bridge functions were estimated using recently developed RKHS-norm regularized adversarial stabilized methods for solving integral equations. A separate theoretical analysis of these estimators—presented in Chapter 4—establishes convergence guarantees under smoothness assumptions on the target solution, without requiring uniqueness. In the present work, we employed these estimators to construct and analyze the bridge-function-based procedure, demonstrating their practical value for causal inference in the presence of unmeasured confounding.

Our numerical experiments demonstrated that the estimator performed as predicted by theory across various scenarios with differing correlations between the negative controls and the unmeasured confounder. As expected, it outperformed conventional non-proximal methods that cannot account for unmeasured confounding. We applied our methodology to the ENSEMBLE trial data, further showcasing its utility.

Future research directions include extending this work to accommodate censored time-to-event outcomes and developing proximal methods for longitudinal modified treatment policies. These extensions would enable the optimization of sequential treatment decisions while addressing time-varying unmeasured confounding.

Chapter 4

SOURCE CONDITION-BASED ANALYSIS AND ASSUMPTION COMPARISONS FOR RKHS-NORM PENALIZED KERNEL ADVERSARIAL STABILIZED ESTIMATORS

4.1 Introduction

In many scientific and policy-relevant applications, target parameters depend on nuisance functions identified through conditional moment restrictions. Estimating such functions often leads to ill-posed inverse problems, where small perturbations in intermediate quantities can cause large estimation errors. These settings arise prominently in causal inference under unmeasured confounding (Santos, 2011; Severini and Tripathi, 2012; Escanciano and Li, 2013; Freyberger and Horowitz, 2015; Miao et al., 2018; Tchetgen et al., 2020; Kallus et al., 2021; Ghassami et al., 2022; Bennett et al., 2022, 2023b; Cui et al., 2024) and in analyses of missing-not-at-random data (Miao and Tchetgen Tchetgen, 2016; Li et al., 2023; Miao et al., 2024).

Classical approaches—such as series or sieve-based Generalized Method of Moments (GMM)—replace the infinite set of implied moment equations with a finite system, obtained by testing against a pre-chosen set of functions. While computationally straightforward, these methods can miss key structural components of the solution when they lie in directions poorly represented by the chosen test space.

Adversarial estimators (Bennett et al., 2019; Muandet et al., 2020b,a; Dikkala et al., 2020; Liao et al., 2020; Kallus et al., 2021; Ghassami et al., 2022; Mastouri et al., 2021; Bennett et al., 2022; Zhang et al., 2023; Bennett et al., 2023a,b; Park et al., 2024) address

this limitation by framing estimation as a two-player zero-sum game: a candidate solution is pitted against an adaptive critic that searches for the most challenging directions of violation. Within this framework, Regularized Adversarial Stabilized (RAS) estimators apply Tikhonov regularization to stabilize estimation. Of particular interest are variants that use reproducing kernel Hilbert spaces (RKHSs) for both estimation and testing. When regularized via the RKHS norm—hereafter referred to as KRAS estimators—they admit closed-form solutions without the numerical instability suffered by their $\mathcal{L}^2(\mathbb{P}_n)$ –penalized RAS counterparts. The latter, while also closed-form, require solving nearly ill-conditioned systems and can be computationally fragile.

Theoretical analysis of $\mathcal{L}^2(\mathbb{P}_n)$ –penalized RAS estimators, hereafter referred to as LRAS estimators, has yielded sharp finite-sample bounds under interpretable source conditions that link solution smoothness to operator smoothing properties. In contrast, existing theory for KRAS estimators is limited: available bounds rely on abstract ill-posedness measures that are hard to interpret or verify and often impose strong conditions such as uniqueness of the solution.

This paper has two goals of equal importance. First, we develop a new theoretical framework for KRAS estimators that yields finite-sample bounds for both weak error and RMSE under transparent source conditions, without requiring uniqueness of the inverse problem solution. These conditions describe the smoothness of the minimal-norm solution relative to the smoothing effect of the composite operator formed by the conditional expectation and the RKHS kernel integral operator. Second, we place KRAS estimators in the broader landscape by systematically comparing their assumptions underpinning their error bounds, and the convergence rates that each deliver, with those of LRAS estimators and with those of recently proposed, computationally stable Kernel Maximal Moment (KMMR) estimators. This comparative perspective situates our results within the spectrum of available methods,

clarifying the trade-offs each entails and the contexts in which their respective assumptions are most plausible.

The rest of the paper is organized as follows. Section 4.2 introduces the general class of target parameters and formulates their identification through conditional moment equations, which often give rise to ill-posed inverse problems. Section B.1.4 reviews the debiased machine learning (DML) framework for estimating the target parameter, emphasizing the importance of controlling both weak error and RMSE in the estimation of nuisance functions. Section 4.4 outlines the broad class of regularized estimation strategies, with a particular focus on adversarial estimators. Section 4.5 presents our main theoretical results: finite-sample bounds on weak error and RMSE of KRAS estimator under interpretable source conditions, without requiring uniqueness of the solution to the inverse problem. Section 4.6 provides comparative analysis of KRAS, LRAS and KMMR estimators. We conclude in Section 4.7 with a discussion of open problems.

4.2 Target Parameter

Let $O_1, \dots, O_n$ be independent copies of a random vector O with unknown probability distribution P_0 assumed to belong to a class $\mathcal{P}$ —to be specified later—of mutually absolutely continuous laws P on some sample space to be defined later. Let $W := (W', X)$ and $Z := (Z', X)$ be two subvectors of O , where X contains the variables common to both W and Z , if any, and W' and Z' contain the remaining, non-overlapping variables. We assume W' and Z' are both non-empty, so that W and Z are distinct but may share the components collected in X .

For any random vector V distributed according to P_V , we let $\mathcal{L}^2(P_V)$ denote the space of real-valued measurable functions on the sample space of V with finite second moment under P_V . We use calligraphic upper case letters to denote the, assumed common, support of

random variables under any law P in $\mathcal{P}$, i.e. if V is a subvector of O , $\mathcal{V}$ denotes the common support of V under any P in $\mathcal{P}$. We denote with $\|\cdot\|_2$ and $\langle \cdot, \cdot \rangle_2$ the norm and the inner product in the $\mathcal{L}^2$ space with respect to the relevant probability law. For any real-valued map $f : \mathcal{V} \mapsto \mathbb{R}$, $\|f\|_\infty := \sup_{v \in \mathcal{V}} |f(v)|$. We let $\mathbb{E}_n$ denote the empirical mean operator.

We are concerned with inference based on $O_1, \dots, O_n$ about the value taken at P_0 of functionals $P \mapsto \theta(P)$ that admit the following triple representation:

$$\theta(P) = \Psi_P(h_{0,P}) = \Phi_P(g_{0,P}) = \mathbb{E}_P \{r(O)h_0(W)g_{0,P}(Z)\} \quad (4.1)$$

where

1. $\Psi_P(h) := \mathbb{E}_P \{\tilde{m}(O; h)\}$, $\Phi_P(g) := \mathbb{E}_P \{m(O; g)\}$ with $\tilde{m} : \mathcal{O} \times \mathcal{L}^2(P_W) \rightarrow \mathbb{R}$ and $m : \mathcal{O} \times \mathcal{L}^2(P_Z) \rightarrow \mathbb{R}$ are known mappings such that for each o , the maps $h \mapsto \tilde{m}(o; h)$ and $g \mapsto m(o; g)$ are linear, and the functionals $h \mapsto \Psi_P(h)$ and $g \mapsto \Phi_P(g)$ are bounded (and linear) on $\mathcal{L}^2(P_W)$ and $\mathcal{L}^2(P_Z)$ respectively, with Riesz representers denoted by $\tilde{\rho}_P(W)$ and $\rho_P(Z)$; and
2. $h_{0,P}$ and $g_{0,P}$ solve respectively the integral equations

$$\mathbb{E}_P \{r(O)h(W)|Z\} = \rho_P(Z) \quad a.e.(P_Z), \quad (4.2)$$

and

$$\mathbb{E}_P \{r(O)g(Z)|W\} = \tilde{\rho}_P(W) \quad a.e.(P_W), \quad (4.3)$$

where r is a known measurable function of O satisfying $\|r\|_\infty < \infty$.

We make no additional assumptions on the data generating process other than the existence of solutions to equations (4.2) and (4.3) at $P = P_0$. Therefore, the aforementioned model $\mathcal{P}$ is the set of all laws P on the sample space of O such that equations (4.2) and (4.3)

have a solution. Importantly, we assume that solutions to these equations exist but we do not require these solutions to be unique. Nevertheless, when both equations admit solutions, the parameter $\theta(P)$ is well-defined, in the sense that its value remains the same regardless of the specific solutions $h_{0,P}$ and $g_{0,P}$ that are used in (4.1). This well known fact (see, e.g. Severini and Tripathi (2012)) can be seen as follows. Let $\mathcal{T} : \mathcal{L}^2(P_W) \rightarrow \mathcal{L}^2(P_Z)$ and $\mathcal{T}^* : \mathcal{L}^2(P_Z) \rightarrow \mathcal{L}^2(P_W)$ denote the operators

$$\mathcal{T}h(Z) := \mathbb{E}_P \{r(O)h(W)|Z\} \quad \text{and} \quad \mathcal{T}^*g(W) := \mathbb{E}_P \{r(O)g(Z)|W\}.$$

The operator $\mathcal{T}_P^*$ is the adjoint operator of $\mathcal{T}_P$. If $h_{0,P}$ and $h_{0,P}^\diamond$ are two solutions of equation (4.2), then $h_{0,P} - h_{0,P}^\diamond$ is in the null space of $\mathcal{T}_P$ and consequently is orthogonal to the range of the adjoint operator $\mathcal{T}_P^*$. On the other hand, the existence of a solution to equation (4.3) implies that $\tilde{\rho}_P(W)$ is in the range of the operator $\mathcal{T}_P^*$, and consequently is orthogonal to $h_{0,P} - h_{0,P}^\diamond$. Therefore, we have $\mathbb{E}_P \{\tilde{m}(O; h_{0,P})\} - \mathbb{E}_P \{\tilde{m}(O; h_{0,P}^\diamond)\} = \mathbb{E}_P \{(h_{0,P} - h_{0,P}^\diamond) \tilde{\rho}_P(W)\} = 0$. The same reasoning can be used to argue the uniqueness of $\mathbb{E}_P \{m(O; g_{0,P})\}$ regardless of the solution $g_{0,P}$ of (4.3).

For some estimands $\theta(P_0)$, see for instance Example 4.2 below, it just happens that it is a-priori known that for some non-empty subset $\mathcal{X}_0 \subsetneq \mathcal{X}$ it holds that $\rho_P(Z) = 0$ a.e. (P_Z) if and only if $Z = (Z', X)$ a.e. (P_Z) and $X \notin \mathcal{X}_0$. In such a case, to enforce this restriction we will restrict attention to solutions of equation (4.2) of the form $h_{0,P}(w) = I_0(x)h_{0,P}^\dagger(w)$ where $h_{0,P}^\dagger$ is any function in $\mathcal{L}^2(P_W)$ that satisfies $\mathcal{T}_P(h)(z) = \rho_P(z)$ whenever $z = (z', x)$ and $x \in \mathcal{X}_0$, where here and throughout

$$I_0(x) := 1 \quad \text{if } x \in \mathcal{X}_0 \quad \text{and} \quad I_0(x) := 0 \quad \text{otherwise.}$$

We follow the same rule for defining the solutions $g_{0,P}$ whenever the Riesz representer $\tilde{\rho}_P(W)$

has structural zeros at some subset of the support of X .

Estimands of the form $\theta(P_0)$ are ubiquitous in causal inference, including settings with unobserved confounding—where identification is achieved by means of instrumental variables (Santos, 2011; Severini and Tripathi, 2012; Escanciano and Li, 2013; Freyberger and Horowitz, 2015) or by means of proxy variables (Miao et al., 2018; Tchetgen et al., 2020; Kallus et al., 2021; Cui et al., 2024)—as well as in missing-not-at-random data scenarios, where shadow variables enable identification (Miao and Tchetgen Tchetgen, 2016; Li et al., 2023; Miao et al., 2024). We review three examples next. Inference for general estimands $\theta(P_0)$, under weak assumptions on h_{0,P_0} and g_{0,P_0} , has been the focus of extensive recent work (Ghassami et al., 2022; Bennett et al., 2022; Chernozhukov et al., 2023; Bennett et al., 2023b). The general approach adopted in these papers, often referred to as debiased machine learning (DML), is reviewed in the next section.

4.2.1 Examples

In this section, we provide several examples of parameters satisfying the representation (4.1).

Example 4.1 (Proximal Inference for Generalized Average Causal Effects). *Suppose that $O = (L, A, Z', W', Y)$, where A is a discrete or continuous treatment, Y an outcome, and L a vector of observed baseline covariates, and Z' and W' are correlates of an unmeasured confounder U , referred to as negative control treatments and negative control outcomes, respectively, satisfying $Z' \perp Y \mid A, L, U$ and $(A, Z') \perp W' \mid L, U$. For each $a \in \mathcal{A}$, let $Y(a)$ denote the potential outcome under treatment level a and $Y(\cdot) := \{Y(a) : a \in \mathcal{A}\}$. Let $O_{full} = (L, A, Z', W', Y, Y(\cdot))$. Let P_0 and $P_{full,0}$ denote the unknown distributions of the factual random vector O and the factual/counterfactual vector O_{full} . Suppose the goal is to estimate a parameter $\gamma(P_{full,0}) := \mathbb{E}_{P_{full,0}} \{Y(A)\pi(A, L)\}$ where $\pi : \mathcal{A} \times \mathcal{L} \mapsto \mathbb{R}$ is a given specified weighting function. In particular, when A is binary and $\pi(a, l) = 2a - 1$, $\gamma(P_{full,0})$ is the aver-*

age treatment effect (ATE), and when $\pi(a|l) = a$, $\gamma(P_{full,0}) = \mathbb{E}_{P_{full,0}}\{Y(1)\}$. Under the consistency assumption $Y(A) = Y$ and additional regularity conditions which essentially ensure that Z' and W' are strong proxies of the unmeasured confounder U , it can be shown (see [Miao et al. \(2018\)](#); [Tchetgen et al. \(2020\)](#); [Kallus et al. \(2021\)](#); [Ghassami et al. \(2022\)](#); [Cui et al. \(2024\)](#)) that $\gamma(P_{full,0})$ is identified and equal to $\theta(P_0)$ of the form (4.1), with $W = (L, A, W')$, $Z = (L, A, Z')$, $r(o) = 1$, $\tilde{m}(o; h) = \pi(a, l) \cdot h(w)$, $m(o; g) = y \cdot g(z)$ and corresponding Riesz representers $\tilde{\rho}_{P_0}(w) = \pi(a, l)/p_{0,A|L,W'}(a|l, w')$ and $\rho_{P_0}(z) = \mathbb{E}_{P_0}(Y \mid Z = z)$, where $p_{0,A|L,W'}$ denotes the conditional density or probability mass function of A given (L, W') . In this example $X := (A, L)$. Note that while ρ_{P_0} does not have structural zeros, $\tilde{\rho}_{P_0}$ will have exactly the same structural zeros as the weight function $\pi(a, l)$. For instance, if A is binary and $\pi(a, l) = a$, then $\tilde{\rho}_{P_0}(w) = 0$ for $w = (a = 0, l, w')$.

Example 4.2 (Proximal Inference for Policy Shift Interventions). Suppose that under the same assumptions as in Example 4.1, with A being continuous, we are interested in estimating the parameter $\gamma(P_{full,0}) := \mathbb{E}_{P_{full,0}}[Y\{q(A)\}]$ where $q : \mathcal{A} \rightarrow \mathcal{A}$ is a known function of the actually received treatment ([Haneuse and Rotnitzky, 2013](#); [Díaz et al., 2023](#)). For instance, q may take the form $q(a) = (a + \delta)I\{a \in [c, d - \delta - \varepsilon]\} + \{a + \delta(d - a)/(\delta + \varepsilon)\}I\{a \in [d - \delta - \varepsilon, d]\}$ where $[c, d]$ is the support of A and δ and ε are fixed pre-specified constants such that $d - \delta - \varepsilon > c$. In Chapter 3, we showed that under certain regularity conditions which include the requirement that q be invertible and differentiable almost everywhere, $\theta(P_{full,0})$ is identified and admits the representation (4.1), with $r(o) = 1$, $\tilde{m}(o; h) = h\{q(a), l, w'\}$, $m(o; g) = y \cdot g(z)$ and corresponding Riesz representers ρ_{P_0} as in Example 4.1—since the map $m(o; g)$ is the same as in that example—and $\tilde{\rho}_{P_0}(w) = I\{a \in [c + \delta, d]\} \cdot \frac{dq^{-1}(a)}{da} \cdot \frac{p_{0,A|L,W}\{q^{-1}(a)|l, w'\}}{p_{0,A|L,W}(a|l, w')}$. Recalling that in this problem $X := (L, A)$, we see that $\tilde{\rho}_{P_0}$ is zero outside the set $\{(l, a) : a \in [c + \delta, d]\}$, so this last set plays the role of $\mathcal{X}_0$ for estimation of the solution of equation (4.3).

Example 4.3 (Missing Not At Random Data with Shadow Variables). Suppose that $O = (\Delta, \Delta Y, X, Z')$, where Δ is binary, Y is an outcome observed if and only if $\Delta = 1$, X is a vector of covariates, and Z' is a variable informative of Y , referred to as shadow variable, satisfying $Z' \perp \Delta \mid X, Y$ and $Z' \not\perp Y \mid X, \Delta = 1$. Let $O_{full} = (\Delta, Y, X, Z)$ denote the full data vector. Let P_0 and $P_{full,0}$ denote the unknown distributions of the observed random vector O and the full data vector O_{full} . Suppose the missingness mechanism is not at random and that the goal is to estimate a parameter $\gamma(P_{full,0}) := \mathbb{E}_{P_{full,0}} \{(1 - \Delta)\pi(X, Y)\}$ where $\pi : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$ is a given function. In particular, when $\pi(X, Y) = Y$, $\gamma(P_{full,0})/P_0(\Delta = 0)$ is the mean outcome among individuals with missing outcomes. Under certain structural assumptions on π ,—namely, that $\pi(X, Y)$ depends only on features of Y that are indirectly captured through Z' given X —it can be shown (Miao and Tchetgen Tchetgen, 2016; Li et al., 2023; Miao et al., 2024) that $\gamma(P_{full,0})$ is identified and equals $\theta(P_0)$ of the form (4.1), with $W = (X, \Delta Y)$, $Z = (X, Z')$, $r(o) = \delta$, $\tilde{m}(o; h) = \delta\pi(x, \delta y)h(x, \delta y)$, $m(o; g) = (1 - \delta)g(x, z)$ and corresponding Riesz representers $\tilde{\rho}_{P_0}(x, \delta y) = \pi(x, \delta y)\mathbb{E}(\Delta \mid X = x, \Delta Y = \delta y)$ and $\rho_{P_0}(x, z) = \mathbb{E}(1 - \Delta \mid X = x, Z = z)$.

4.3 Debiased Estimation of $\theta(P_0)$

In this section we review the general debiased-machine learning estimation strategy for estimating parameters of the form $\theta(P_0)$ as defined in (4.1). It is well known that if an estimator $\hat{h}$ is obtained using flexible statistical methods, the plug-in estimator $\hat{\Psi}(\hat{h}) := n^{-1} \sum_{i=1}^n m(O_i, \hat{h})$ does not generally achieve $\sqrt{n}$ -rate convergence to $\theta(P_0)$ even when $\hat{h}$ converges in $\mathcal{L}^2(P_{0,W})$ to a solution h_{0,P_0} . The main reason is that the dependence of

$$\Psi_{P_0}(\hat{h}) - \Psi_{P_0}(h_{0,P_0}) = \int (\hat{h} - h_{0,P_0})(w) \tilde{\rho}_{P_0}(w) dP_0(o)$$

on $\hat{h} - h_{0,P_0}$ causes the plug-in estimator to inherit the bias of $\hat{h}$. Writing

$$\begin{aligned} \int (\hat{h} - h_{0,P_0})(w) \tilde{\rho}_{P_0}(w) dP_0(o) &= \int \hat{h}(w) \tilde{\rho}_{P_0}(w) dP_0(o) - \int m(o, g_{0,P_0}) dP_0(o) \\ &= \int r(o) \hat{h}(w) g_{0,P_0}(z) dP(o) - \int m(o, g_{0,P_0}) dP_0(o) \end{aligned}$$

suggests estimating the bias of the plug-in estimator with $n^{-1} \sum_{i=1}^n \{r(O_i) \hat{h}(W_i) \hat{g}(Z_i) - m(O_i, \hat{g})\}$, where $\hat{g}$ is an estimator converging to a solution $g_{0,P}$ constructed using flexible statistical methods. Subtracting this estimated bias from the plug-in estimator, yields the bias-corrected estimator of $\theta(P_0)$ given by $n^{-1} \sum_{i=1}^n \phi(O_i; \hat{h}, \hat{g})$, where for any h and g ,

$$\phi(O; h, g) := \tilde{m}(O; h) + m(O; g) - r(O)h(W)g(Z).$$

The map $o \mapsto \phi(o; h_{P_0}, g_{P_0}) - \theta(P_0)$ corresponds to an influence function at P_0 of the functional $P \mapsto \theta(P)$ in model $\mathcal{P}$ (see, e.g., [Kennedy \(2024\)](#) and references therein).

The preceding debiased estimator may still fail to be asymptotically linear when flexible machine learning methods are employed to compute $\hat{h}$ and $\hat{g}$ due to overfitting. A popular approach to addressing this issue is to integrate debiasing with a technique called cross-fitting. Specifically, the sample is partitioned into K roughly equal sized subsamples, the data in all but one sample, say the k^{th} sample, is used to compute the estimators of the nuisance functions, denoted as $\hat{h}^{(-k)}$ and $\hat{g}^{(-k)}$, and the estimator of $\theta(P_0)$ is computed as $\hat{\theta} = n^{-1} \sum_{k=1}^K \sum_{i \text{ in sample } k} \phi(O_i; \hat{h}^{(-k)}, \hat{g}^{(-k)})$. This estimator belongs to the class of DML estimators ([Chernozhukov et al., 2018](#)), and it is also commonly described as an influence-function-based one-step estimator with cross-fitting ([Robins et al., 2008](#)).

The following theorem establishes the asymptotic behavior of $\hat{\theta}$, a result that is now standard in the literature on semiparametric and DML estimation, and whose proof we consequently omit (see [Robins et al. \(2008\)](#); [Chernozhukov et al. \(2018, 2023\)](#) and [Kennedy](#)

(2024)).

Theorem 4.1. *Suppose that h_{0,P_0} and g_{0,P_0} solve equations (4.2) and (4.3), respectively, at $P = P_0$. Suppose that $\int \{\phi(o; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(o; h_{0,P_0}, g_{0,P_0})\}^2 dP_0(o) = o_P(1)$ holds for each $k \in \{1, \dots, K\}$. Then, $\sqrt{n}\{\hat{\theta} - \theta(P_0)\} = \sqrt{n}\mathbb{E}_n[\phi(O; h_{0,P_0}, g_{0,P_0}) - \theta(P_0)] + \sqrt{n} R_n + o_P(1)$, where*

$$R_n := \sum_{k=1}^K \frac{n_k}{n} \int r(o) \left\{ \hat{h}^{(-k)}(w) - h_{0,P_0}(w) \right\} \left\{ g_{0,P_0}(z) - \hat{g}^{(-k)}(z) \right\} dP_0(o),$$

with n_k denoting the size of the k^{th} subsample. In particular, if $\mathbb{E}_{P_0}\{\phi(O; h_{0,P_0}, g_{0,P_0})^2\} < \infty$ and $R_n = o_P(n^{-1/2})$, then $\sqrt{n}\{\hat{\theta} - \theta(P_0)\} \xrightarrow{d} \mathcal{N}(0, \tau_0^2)$, where $\tau_0^2 = \mathbb{E}_{P_0}[\{\phi(O; h_{0,P_0}, g_{0,P_0}) - \theta(P_0)\}^2]$.

Theorem 4.1 highlights the importance of constructing estimators $\hat{h}^{(-k)}$ and $\hat{g}^{(-k)}$ that converge sufficiently quickly to solutions h_{0,P_0} and g_{0,P_0} of the integral equations (4.2) and (4.3), respectively, in order to ensure that $R_n = o_P(n^{-1/2})$. As noted in Bennett et al. (2022) and Chernozhukov et al. (2023), a sufficient condition on the convergence rates of the nuisance estimators to ensure that $R_n = o_P(n^{-1/2})$ can be derived from the following derivation:

$$\begin{aligned} |R_n| &\leq \sum_{k=1}^K \frac{n_k}{n} \left| \int \left[\int r(o) \left\{ \hat{h}^{(-k)}(w) - h_{0,P_0}(w) \right\} dP_0(o|z) \right] \left\{ g_{0,P_0}(z) - \hat{g}^{(-k)}(z) \right\} dP_0(z) \right| \\ &\leq \sum_{k=1}^K \frac{n_k}{n} \left\| \mathcal{T}_{P_0}(\hat{h}^{(-k)} - h_{0,P_0}) \right\|_2 \left\| \hat{g}^{(-k)} - g_{0,P_0} \right\|_2 \end{aligned}$$

where the inequality follows from the Cauchy–Schwarz inequality. A symmetric bound can

be obtained by integrating first with respect to $dP_0(o|w)$. Combining both bounds, yields:

$$|R_n| \leq \sum_{k=1}^n \frac{n_k}{n} \min\{\|\mathcal{T}_{P_0}(\hat{h}^{(-k)} - h_{0,P_0})\|_2 \|\hat{g}^{(-k)} - g_{0,P_0}\|_2, \|\mathcal{T}_{P_0}^*(\hat{g}^{(-k)} - g_{0,P_0})\|_2 \|\hat{h}^{(-k)} - h_{0,P_0}\|_2\}$$

Thus, we upper bound R_n with the minimum of two products: the weak error of one estimator times the RMSE of the other—by weak error, say of estimator $\hat{h}^{(k)}$, we mean $\|\mathcal{T}_{P_0}(\hat{h}^{(k)} - h_{0,P_0})\|_2$. We can also upper bound R_n using the looser inequality:

$$|R_n| \leq \|r\|_\infty \sum_{k=1}^n \frac{n_k}{n} \left\| \hat{g}^{(-k)} - g_{0,P_0} \right\|_2 \left\| \hat{h}^{(-k)} - h_{0,P_0} \right\|_2$$

This later bound is less sharp than the former, since weak errors tend to be substantially smaller than RMSEs. This observation highlights a limitation of the existing analyses of KMMR estimators because they provide only bounds for RMSE: relying exclusively on RMSE bounds can lead to overly stringent smoothness conditions when establishing $\sqrt{n}$ -consistency of the DML estimator of $\theta(P_0)$.

4.4 Adversarial Estimation of the Nuisance Functions

For the remainder of the paper, we focus on estimating the nuisance function that solves equation (4.2). The same analysis applies to equation (4.3) after an appropriate redefinition of the relevant objects. Furthermore, to simplify notation, throughout we will omit the subscript P from the operator $\mathcal{T}_P$, the solution $h_{0,P}$, the Riesz representer ρ_P , and all other operators and quantities, e.g. expectations, defined throughout that depend on P .

Estimating solutions to integral equations is inherently challenging due to their ill-posed nature: small perturbations in the input function ρ can lead to large deviations in the solution. A naive way to circumvent this instability is to restrict attention to a sufficiently small subspace $\mathcal{H}$ of $\mathcal{L}^2(P_W)$ in which the inverse problem becomes well-posed. When $\mathcal{H}$

is finite-dimensional, the resulting estimator is computationally tractable, but the approach is theoretically unsatisfying as it relies on the strong assumption that a solution lies in the small subspace $\mathcal{H}$.

A broadly applicable alternative strategy is to regularize the estimation problem. This can be done by either allowing the finite-dimensional space $\mathcal{H}$ to grow with the sample size, as in sieve-based methods, or by penalizing the complexity of functions in an infinite-dimensional space $\mathcal{H}$, such as an RKHS. Both strategies yield estimators that approximate the solution of the original ill-posed problem as the level of regularization is gradually reduced (Carrasco et al., 2007; Cavalier, 2011). From a unified perspective, several regularization approaches can be viewed as targeting solutions to surrogate regularized optimization problems of the form

$$\arg \min_{h \in \mathcal{H}} \{ \|\rho - \mathcal{T}I_0 h\|^2 + \lambda_{\mathcal{H}} \|h - h^{\diamond}\|^2 \}, \quad (4.4)$$

where $h^{\diamond}$ is a fixed reference function, typically set to zero. These estimators differ in their choice of function class $\mathcal{H}$ (typically a subset of a normed space), in whether they include a penalization term (i.e., whether $\lambda_{\mathcal{H}}$ is positive), and in the choice of norms used for both the loss term $\|\rho - \mathcal{T}I_0 h\|^2$ and the regularization term $\|h - h^{\diamond}\|^2$.

When $\lambda_{\mathcal{H}} = 0$, a common choice is to take $\mathcal{H}$ as a finite-dimensional subspace spanned by the first p elements of an orthonormal basis of $\mathcal{L}^2(P_W)$ and to use the $\mathcal{L}^2(P_W)$ -norm for the loss. In this setting, regularization is relaxed by allowing the dimension p to grow with the sample size. Estimators constructed in this way are referred to as sieve-based and include methods proposed in Newey and Powell (2003) and Ai and Chen (2003, 2007). In contrast, when $\lambda_{\mathcal{H}} \neq 0$, the approach is referred to as Tikhonov regularization. An influential earlier example of this approach is the work of Chen and Pouzo (2012), who consider nonparametric conditional moment models and apply Tikhonov regularization by penalizing the $\mathcal{L}^2$ -norm

of the solution in a sieve space.

In more recent proposals, $\mathcal{H}$ is typically taken to be an infinite-dimensional function class, often restricted via a norm constraint, or a very high-dimensional normed space, like shape constrained functions, neural networks (NN), RKHSs, and random forests. The penalty parameter $\lambda_{\mathcal{H}}$ is strictly positive and regularization is relaxed by letting it decrease to 0. For instance, [Dikkala et al. \(2020\)](#) allow $\mathcal{H}$ to be any of the aforementioned spaces; [Ghassami et al. \(2022\)](#) focus on RKHSs; and [Bennett et al. \(2022\)](#) consider both RKHSs and NNs. In all of these papers the loss is measured with the $\mathcal{L}^2$ -norm, while the penalty is defined via the norm of the function space $\mathcal{H}$ with $h^{\diamond} = 0$. By contrast, other recent approaches use the $\mathcal{L}^2$ -norm for both the loss and the penalty terms. This includes [Liao et al. \(2020\)](#), who choose $\mathcal{H}$ as NNs; and [Bennett et al. \(2023b\)](#), who consider norm-constrained high-dimensional linear models, RKHSs, and NNs. The latter also considers an iterated Tikhonov approach, starting with an initial estimator targeting the solution $h_{\lambda}^{(0)}$ to the regularized problem with $h^{\diamond} = 0$. At each subsequent step, the estimator targets the solution $h_{\lambda}^{(k)}$ to the same problem but with $h^{\diamond}$ set to the previous iterate $h_{\lambda}^{(k-1)}$. This process is repeated for a fixed number of iterations or an adaptively chosen stopping point.

In the Maximal Moment Restriction (MMR) proposals of [Muandet et al. \(2020a\)](#); [Mastouri et al. \(2021\)](#); [Kallus et al. \(2021\)](#); [Zhang et al. \(2023\)](#); and [Park et al. \(2024\)](#), the surrogate problem takes the form:

$$\arg \min_{h \in \mathcal{H}} \max_{g \in \mathcal{G}} [\mathbb{E}\{m(O; g) - h(W)g(Z)\}]^2 + \lambda_{\mathcal{H}} \|h\|^2.$$

We discuss the MMR proposals in Section [4.6.2](#).

In the next subsection, we review the class of RAS estimators. To avoid notational burden, we will drop the superscript $(-k)$ from $\hat{h}^{(-k)}$ and the subscript k from n_k . Furthermore,

we will use $\mathbb{E}_n$ to denote the empirical mean operator calculated with data in the entire sample but the k^{th} sample.

4.4.1 Motivation and Rationale for Regularized Adversarial Stabilized Estimation Strategies

In addition to the statistical advantages discussed in the introduction—most notably, their ability to adaptively focus on directions where the conditional moment restriction is most difficult to satisfy—adversarial estimation strategies offer two key computational benefits. First, they bypass the need to estimate the conditional mean operator $\mathcal{T}$. Second, they do not require prior knowledge of the functional form of the Riesz representer ρ , which facilitates greater automation of the estimation procedure. Methods that do not require knowledge of the functional form of the Riesz representer are referred to in the literature as auto-debiased machine learning (auto-DML). They were first proposed in the case where $X = W = Z$ by Chernozhukov et al. (2022a,b,c), and later extended to the setting where W and Z are distinct subvectors of O by Bennett et al. (2022, 2023b).

The RAS estimation strategies we introduce next, unlike the alternative MMR class of adversarial estimators studied in Muandet et al. (2020a); Zhang et al. (2023); Mastouri et al. (2021); and Park et al. (2024), have the advantage that existing analytical results provide bounds for the weak error. As discussed in Section B.1.4, this makes it possible to establish root- n consistency of the resulting estimator of $\theta(P_0)$ under weaker conditions on the solutions to equations (4.2) and (4.3).

Quite generally, RAS estimators are of the form $\hat{h}(W) = I_0(X)\tilde{h}(W)$ where

$$\tilde{h} := \arg \min_{h \in \mathcal{H}} \max_{g \in \mathcal{G}} [\mathbb{E}_n \{m(O; g) - I_0(X)r(O)h(W)g(Z) - c^2 I_0(X)g(Z)^2\} - \lambda_{\mathcal{G}} \|g\|^2 + \lambda_{\mathcal{H}} \|h\|^2] \quad (4.5)$$

for some constant $c \neq 0$ and, $\mathcal{G}$ and $\mathcal{H}$ specified function classes (Muandet et al., 2020b;

Dikkala et al., 2020; Liao et al., 2020; Kallus et al., 2021; Ghassami et al., 2022; Bennett et al., 2022, 2023b). As noted earlier, Bennett et al. (2023b) also consider iterated Tikhonov estimators with penalties centered at the estimator from the previous iteration. The aforementioned proposals differ on the choice of norms $\|\cdot\|$ for g and h . Additionally, some proposals set $\lambda_{\mathcal{H}} = 0$ and/or $\lambda_{\mathcal{G}} = 0$. See Appendix C.3.1.

To motivate the RAS approach we first observe that for any $g \in L^2(P_Z)$, it holds that $\mathbb{E}[m(O; g)] = \mathbb{E}[\rho(Z)g(Z)] = \mathbb{E}[r(O)h(W)g(Z)]$, where the first equality is because ρ is the Riesz representer, and the second is by h_0 satisfying equation (4.2). In particular, this implies that for any $g \in \mathcal{L}^2(P_Z)$, the residual

$$m(O; g) - I_0(X)r(O)g(Z)h(W) \quad (4.6)$$

has mean equal to 0 at $h = h_0^\dagger$ where $h_0 = I_0 h_0^\dagger$ is a solution of equation (4.2). Consequently,

$$Q_{\mathbb{J}}(h) := \sum_{j=1}^{\mathbb{J}} \{\mathbb{E}[m(O; g_j) - I_0(X)r(O)g_j(Z)h(W)]\}^2$$

for given functions $g_j, j = 1, \dots, J$ where $\mathbb{J} \subseteq \mathbb{N}$, attains its minimum at $h_0^\dagger$. One can then use the empirical version of $Q_{\mathbb{J}}(h)$ as the basis for deriving estimators of $h_0^\dagger$ —a strategy that underlies the classical GMM approach. The functions g_j are known in the machine learning literature as *test functions*. The question then is: how many and which test functions do we choose? The following result sheds light into the answer. In what follows, given a subspace $\mathcal{G}$ of $\mathcal{L}^2(P_Z)$, we let $I_0 \mathcal{G} := \{I_0 g : g \in \mathcal{G}\}$ and we let $\overline{I_0 \mathcal{G}}$ denote its closure in $\mathcal{L}^2(P_Z)$.

Proposition 4.1. *Suppose $\{g_j = I_0 g_j^\dagger : g_j^\dagger \in \mathcal{G}, j \in \mathbb{J}\}$ is an $\mathcal{L}^2(P_Z)$ -orthonormal basis of $\overline{I_0 \mathcal{G}}$. Let $\mathcal{P}_{\overline{I_0 \mathcal{G}}} : \mathcal{L}^2(P_Z) \rightarrow \overline{I_0 \mathcal{G}}$ denote the projection operator. Then $Q_{\mathbb{J}}(h) = \|\mathcal{P}_{\overline{I_0 \mathcal{G}}} \mathcal{T} I_0(h_0^\dagger - h)\|_2^2$.*

To understand the relevance of this result avoiding distracting heavy notation, suppose for the purpose of this discussion that $I_0 = 1$ and that the integral equation (4.2) admits a unique solution h_0 lying in a known function class $\mathcal{H}$. Proposition 4.1 implies that estimators minimizing $Q_{\mathbb{J}}(h)$ over $\mathcal{H}$ implicitly target solutions to the projected equation

$$\mathcal{P}_{\overline{\mathcal{G}}} \circ \mathcal{T} h = \mathcal{P}_{\overline{\mathcal{G}}} \rho, \quad (4.7)$$

where $\overline{\mathcal{G}}$ denotes the closure of the linear span of the test functions $\{g_j\}_{j \in \mathbb{J}}$. If $\overline{\mathcal{G}}$ is a strict subspace of $\mathcal{L}^2(P_Z)$, then any function of the form $h_0 + h^\diamond$, with $\mathcal{T}h^\diamond \in \text{Null}(\mathcal{P}_{\overline{\mathcal{G}}})$, also solves (4.7). Consequently, minimizing $Q_{\mathbb{J}}$ does not distinguish h_0 from these alternatives.

This ambiguity cannot be resolved by selecting the minimum norm solution of (4.7), as this solution will generally differ from h_0 , the unique solution of the original problem (4.2). To see this, suppose that W and Z have no overlap (i.e., X is null), in which case $\mathcal{T} : \mathcal{L}^2(P_W) \rightarrow \mathcal{L}^2(P_Z)$ is compact. Let $(\mu_i, \varphi_i, \psi_i)$ for $i = 1, 2, \dots$ denote the singular value decomposition of $\mathcal{T}$. Since, by assumption, equation (4.2) has a solution, then the Riesz representer ρ lies in the range of $\mathcal{T}$, and can therefore be written as $\rho = \sum_{i: \mu_i \neq 0} \langle \rho, \psi_i \rangle_2 \psi_i$. The minimum-norm solution to the projected equation (4.7) is then $h_0^{\text{proj}} = \sum_{i: \mu_i \neq 0, \psi_i \in \overline{\mathcal{G}}} \mu_i^{-1} \langle \rho, \psi_i \rangle_2 \varphi_i$, while the solution of the equation of interest (4.2) is $h_0 = \sum_{i: \mu_i \neq 0} \mu_i^{-1} \langle \rho, \psi_i \rangle_2 \varphi_i$. If $\rho \notin \overline{\mathcal{G}}$, then there exists some i for which $\langle \rho, \psi_i \rangle_2 \neq 0$ and $\psi_i \notin \overline{\mathcal{G}}$, meaning that h_0^{proj} omits the corresponding direction φ_i . In the machine learning literature, this is referred to as *misspecification of the test class* $\mathcal{G}$. In particular, if ρ has nonzero projection onto infinitely many directions ψ_i , then any finite-dimensional test class $\mathcal{G}$ will necessarily be misspecified. This is the case in sieve-based GMM estimation, where the dimension of $\mathcal{G}$ increases with sample size, and the test functions $\{g_j\}$ are chosen to minimize approximation error. However, for any fixed sample size, the sieve space cannot capture all directions—especially the high-frequency components

of ρ —and as a result, the corresponding estimator will fail to recover those components of the true solution h_0 .

To avoid misspecification, the test class $\mathcal{G}$ must be dense in $\mathcal{L}^2(P_Z)$. However, when Z has continuous components, this requires indexing the class by the entire natural numbers, i.e., $\mathbb{J} = \mathbb{N}$, implying that the objective function $Q_{\mathbb{J}}$ involves an infinite sum. As a result, the empirical version of this criterion becomes computationally infeasible, since one cannot evaluate or minimize an infinite sum in practice. Remarkably, when $\mathcal{T}(h_0 - h)$ lies in $\mathcal{G}$ there exists an equivalent reformulation that avoids direct evaluation of the infinite sum. Specifically, let c be any non-zero constant and define

$$M_P(h) := 4c^2 \max_{g \in \mathcal{G}} \mathbb{E}_P \{m(O; g) - I_0(X)r(O)h(W)g(Z) - c^2 I_0(X)g(Z)^2\}.$$

Proposition 4.2. *For any constant $c \neq 0$, any $h_0^\dagger$ such that $h_0 = I_0 h_0^\dagger$ is a solution of equation (4.2), and any $h \in \mathcal{L}^2(P_W)$, it holds that:*

1. *If $\mathcal{G} \subset \mathcal{L}^2(P_Z)$ such that $(2c^2)^{-1} \mathcal{T}(h_0^\dagger - h) \in \mathcal{G}$ then $M_P(h) = \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2$.*
2. *Suppose $\mathcal{G}$ is a subspace of $\mathcal{L}^2(P_Z)$. Let $I_0 \mathcal{G} := \{I_0 g : g \in \mathcal{G}\}$ and let $\overline{I_0 \mathcal{G}}$ denote its closure in $\mathcal{L}^2(P_Z)$. Let $\{g_j = I_0 g_j^\dagger : g_j^\dagger \in \mathcal{G}, j \in \mathbb{J}\}$ denote an $\mathcal{L}^2(P_Z)$ –orthonormal basis of $\overline{I_0 \mathcal{G}}$ where $\mathbb{J} \subseteq \mathbb{N}$. If $\mathcal{P}_{\overline{I_0 \mathcal{G}}} \circ \mathcal{T}_P I_0(h_0^\dagger - h) \in I_0 \mathcal{G}$, then $Q_{\mathbb{J}}(h) = M_P(h)$.*

Proposition 4.2 is ultimately a consequence of the Fenchel conjugate formulation of the square function (Borwein and Lewis, 2006). Part 1 of this proposition has been shown by Dikkala et al. (2020) and Kallus et al. (2021). Part 2 of the proposition has been established in the finite-dimensional setting by Bennett et al. (2019) and Muandet et al. (2020b), for the specific case $m(O; g) = Yg(Z)$, with Y and Z denoting distinct components of O . To the best of our knowledge, the extension of this result to infinite-dimensional function classes has not yet been reported in the literature. Proposition 4.2 has the following key implication.

Corollary 4.1. *Let $\mathcal{G} \subset \mathcal{L}^2(P_Z)$ and let $c \neq 0$ be a constant. If there exists $h_0^\dagger$ such that $h_0 = I_0 h_0^\dagger$ solves equation (4.2), and for any $h \in \mathcal{H}$ it holds that $(2c^2)^{-1} \mathcal{T}_P(h_0^\dagger - h) \in \mathcal{G}$, then*

$$\begin{aligned} & \arg \min_{h \in \mathcal{H}} \{ \|\rho - \mathcal{T} I_0 h\|_2^2 + \lambda_{\mathcal{H}} \|h\|^2 \} \\ &= \arg \min_{h \in \mathcal{H}} \left[\max_{g \in \mathcal{G}} \mathbb{E}_P \{ m(O; g) - I_0(X) r(O) h(W) g(Z) - c^2 I_0(X) g(Z)^2 \} + \lambda_{\mathcal{H}} \|h\|^2 \right]. \end{aligned}$$

It is now clear from Corollary 4.1 that when $\mathcal{G}$ is a linear space, under the *closeness* condition

$$\mathcal{T}(h_0^\dagger - h) \in \mathcal{G} \quad \text{holds for every } h \in \mathcal{H}, \quad (4.8)$$

the RAS estimators (4.5) with $\lambda_{\mathcal{G}} = 0$ are estimating solutions of surrogate problems of the form (4.4) that employ the $\mathcal{L}^2(P_Z)$ -norm for the loss term. If $\mathcal{H}$ is a normed space and it is known that $\|h\|_{\mathcal{H}} < B$ for some constant B , then the outer minimization can be restricted to functions in $\mathcal{H}$ with norm at most B . In this case, the closeness condition (4.8) only needs to hold over that restricted subset.

On the other hand, when the closeness condition (4.8) does not hold, part 2 of Proposition 4.2 tells us that the RAS estimators can still be interpreted as regularized GMM estimators over a class of test functions that are dense in the $\mathcal{L}^2(P_Z)$ -closure of the class $I_0 \mathcal{G}$. The requirement that the projection of $\mathcal{T}_P I_0(h_0^\dagger - h)$ onto $\overline{I_0 \mathcal{G}}$ lies in $I_0 \mathcal{G}$ is equivalent to requiring that the maximum over $\mathcal{G}$ of the expected stabilized residuals exists. Existing finite-sample error bounds for RAS estimators have been established only under the stronger closeness condition (4.8), but the preceding discussion on test-class misspecification suggests a potential relaxation: if one is willing to assume that $\rho \in \mathcal{G}$, then it may be possible to derive analogous bounds under the weaker condition that the projection of $\mathcal{T} I_0(h_0^\dagger - h)$ onto $\overline{I_0 \mathcal{G}}$ belongs to $I_0 \mathcal{G}$.

KRAS estimators $\tilde{h}$ of the form (4.5) with $\lambda_{\mathcal{H}} > 0$, $\lambda_{\mathcal{G}} > 0$, $\mathcal{H}$ and $\mathcal{G}$ taken as RKHSs, and penalization based on the corresponding RKHS norms were first studied by Dikkala et al. (2020) in the special case $m(o; g) = y \cdot g(z)$, and by Ghassami et al. (2022) in the more general setting $m(o; g) = s(o)g(z)$, where $s(\cdot)$ is a known function satisfying $\|s\|_{\infty} < \infty$. In both works, the set $\mathcal{X}_0$ was taken to be the entire domain $\mathcal{X}$, equivalently assuming $I_0 \equiv 1$. Choosing $\mathcal{H}$ and $\mathcal{G}$ as RKHSs and penalizing via their respective RKHS norms is attractive because it yields computationally stable estimators that admit a closed-form expression. In contrast, omitting the $\lambda_{\mathcal{G}}\|g\|_{\mathcal{G}}^2$ penalty or replacing RKHS norms with empirical $L^2(\mathbb{P}_n)$ norms for the penalization of h may still result in closed-form estimators, but their computation typically involves inversion of nearly ill-posed $n \times n$ matrices. Including the $\lambda_{\mathcal{G}}\|g\|_{\mathcal{G}}^2$ penalization and using RKHS norms for penalizing h leads instead to inversion of well-posed matrices, thus stabilizing the computation. See Appendix B.2.1 for the closed form expression of KRAS estimators.

Dikkala et al. (2020) and Ghassami et al. (2022) provided a finite-sample analysis of KRAS estimators. Their analyses begin by showing that, under suitable regularity conditions and with appropriately chosen tuning parameters $\lambda_{\mathcal{G}}$ and $\lambda_{\mathcal{H}}$, the weak error $\|\mathcal{T}(\hat{h} - h_0)\|_2$ is of order $O(\delta_n)$, where δ_n is an upper bound on the critical radius of the classes $\mathcal{H}$ and $\mathcal{G}$, and h_0 is any solution to equation (4.2). They then assume that the solution h_0 to equation (4.2) is unique and derive a bound on the RMSE of the form

$$\|\hat{h} - h_0\|_2 \leq \tau(\delta_n), \quad (4.9)$$

where, for any $\delta > 0$, $\tau(\delta) := \sup \{\|h\|_2 : h \in \mathcal{H}, \|h\|_{\mathcal{H}} \leq 2B, \|\mathcal{T}h\|_2 \leq \delta\}$ with B being a universal constant such that $\|\hat{h}\|_{\mathcal{H}} \leq B$ and $\|h_0\|_{\mathcal{H}} \leq B$. Note that $\tau(\delta)$ is a measure of the ill-posedness of equation (4.2) that depends both on the data generating process and on $\mathcal{H}$.

Unfortunately, if there exists a nonzero function $h^{(0)} \in \mathcal{N}(\mathcal{T}) \cap \mathcal{H}$, where $\mathcal{N}(\mathcal{T})$ denotes the null set of the conditional mean operator $\mathcal{T}$, then the bound (4.9) becomes vacuous; that is, $\tau(\delta_n) > C > 0$ for some universal constant C . Indeed, $h^{(0)*} := B h^{(0)} / \|h^{(0)}\|_{\mathcal{H}}$ satisfies $\mathcal{T}h^{(0)*} = 0$ and $\|h^{(0)*}\|_{\mathcal{H}} = B < 2B$, implying that $\tau(\delta) \geq \|h^{(0)*}\|_2 =: C$ for all $\delta > 0$. Similar comments were made by Kallus et al. (2021) and Bennett et al. (2023b).

Even under the condition $\mathcal{N}(\mathcal{T}) \cap \mathcal{H} = 0$, there is no known general methodology for expressing or bounding $\tau(\delta)$ as a function of δ . Dikkala et al. (2020) obtained such a bound in a particular setting, relying on a strong structural assumption on the data-generating process that, in effect, requires Z to be a strong predictor of W . While this result illustrates that bounding $\tau(\delta)$ is possible in certain special cases, to our knowledge no general approach is currently available.

In contrast, more satisfactory finite-sample analyses exist for the computationally unstable RAS estimators that employ $\mathcal{H}$ and $\mathcal{G}$ as RKHSs with $\mathcal{L}^2(\mathbb{P}_n)$ penalties (Liao et al., 2020; Bennett et al., 2023b). These analyses adopt a different notion of ill-posedness that does not require uniqueness of the solution to the integral equation, focusing instead on the minimum $\mathcal{L}^2(P_W)$ -norm solution satisfying the so-called *β -source condition*.

To review the definition of the β -source condition and its implication for controlling the regularization bias in complete generality, consider an arbitrary bounded linear operator $\mathcal{A} : \mathcal{E}_1 \rightarrow \mathcal{E}_2$, where $\mathcal{E}_1$ and $\mathcal{E}_2$ are Hilbert spaces, and suppose we aim to solve the equation $\mathcal{A}h = \rho$ for some $\rho \in \mathcal{E}_2$. Assume a solution to the equation exists, and let h_0 denote the solution with minimal $\mathcal{E}_1$ -norm. Let $\mathcal{A}^* : \mathcal{E}_2 \rightarrow \mathcal{E}_1$ denote the adjoint of the operator $\mathcal{A}$. The β -source condition posits that h_0 lies in the range of the fractional power operator $(\mathcal{A}^*\mathcal{A})^{\beta/2}$; that is,

$$h_0 = (\mathcal{A}^*\mathcal{A})^{\beta/2}h^* \quad \text{for some } h^* \in \mathcal{E}_1. \quad (4.10)$$

When $\mathcal{A}$ is a compact operator, this condition amounts to assuming that h_0 concentrates

its energy in the well-identified directions of the inverse problem—i.e., the directions along which the operator $\mathcal{A}^*\mathcal{A}$ is most stable (see Chapter 3 of Carrasco et al. (2007) for the interpretation of the source condition when $\mathcal{A}$ is compact and Chapter 4 of Hohage (2002) for the definition of the power of a bounded operator). By excluding the possibility that h_0 has substantial components in directions with near-zero singular values of $\mathcal{A}$, the assumption facilitates control over $\|h_\lambda - h_0\|_{\mathcal{E}_1}$ where

$$h_\lambda := \arg \min_{h \in \mathcal{E}_1} \|\mathcal{A}h - \rho\|_{\mathcal{E}_2}^2 + \lambda \|h\|_{\mathcal{E}_1}^2 \quad (4.11)$$

is the solution to the surrogate Tikhonov-regularized problem. The error bounds are given by the following classical result from the theory of inverse problems (see Theorem 5.2 of Hohage (2002)).

Theorem 4.2. *Let $\mathcal{A} : \mathcal{E}_1 \rightarrow \mathcal{E}_2$ be a bounded linear operator between the Hilbert spaces $\mathcal{E}_1$ and $\mathcal{E}_2$ and let $\rho \in \mathcal{E}_2$ be in the range of $\mathcal{A}$. Suppose the solution h_0 to the equation $\mathcal{A}h = \rho$ with minimal $\mathcal{E}_1$ -norm satisfies the β -source condition (4.10). Let h_λ be defined as in (4.11). Then, there exist constants c_1 and c_2 such that $\|h_0 - h_\lambda\|_{\mathcal{E}_1}^2 < c_1 \lambda^{\min\{\beta, 2\}}$ and $\|\mathcal{A}(h_0 - h_\lambda)\|_{\mathcal{E}_2}^2 < c_2 \lambda^{\min\{\beta+1, 2\}}$.*

In the next section, we provide an alternative finite-sample analysis for estimators of the form $\hat{h} = I_0 \tilde{h}$, where $\tilde{h}$ is a KRAS estimator. We derive informative finite-sample bounds on both the weak error and the root mean squared error, under two variants of the β -source condition imposed on the solution with minimal $\mathcal{H}$ -norm. Crucially, we do not assume $\mathcal{N}(\mathcal{T}) \cap \mathcal{H} = 0$, nor do we rely

In Section 4.6 we argue that while adversarial stabilized estimators with $\mathcal{L}^2(\mathbb{P}_n)$ penalties can achieve faster convergence rates under comparable smoothness assumptions, these rates rely on a strong condition that KRAS does not require—and may fail entirely when that

condition is violated. We conjecture that, under these assumptions, the slower rates of KRAS are not a limitation of our analysis—i.e., they do not stem from overly loose bounds—but rather reflect an intrinsic feature of the KRAS methodology itself: by penalizing with the RKHS norm instead of the $\mathcal{L}^2(\mathbb{P}_n)$ norm, the KRAS estimator implicitly tackles a more severely ill-posed inverse problem. Thus, KRAS trades statistical speed for computational stability and less stringent assumptions.

4.5 Finite-sample Bounds on the Errors of the KRAS Estimators

In this section we will focus on the KRAS estimator $\hat{h}_{KRAS}(W) := I_0(X)\tilde{h}_{KRAS}(W)$, where

$$\begin{aligned} \tilde{h}_{KRAS} := \arg \min_{h \in \mathcal{H}_B} \max_{g \in \mathcal{G}} & \left[\mathbb{E}_n \{ m(O; g) - I_0(X)r(O)h(W)g(Z) - c^2 I_0(X)g(Z)^2 \} \right. \\ & \left. - \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2 \right], \end{aligned} \quad (4.12)$$

$\mathcal{H}$ and $\mathcal{G}$ are RKHSs with kernels $K_{\mathcal{H}}$ and $K_{\mathcal{G}}$, and $\mathcal{H}_B := \{h \in \mathcal{H} \mid \|h\|_{\mathcal{H}} \leq B\}$ with B a constant defined later after Assumption 4.2.

We shall introduce two distinct β -source conditions and derive finite-sample bounds for both the RMSE and the weak error under each condition and additional regularity conditions. The source conditions differ in the operators that define them. The first arises directly from an application of Theorem 4.2. The second is obtained after a transformation that clarifies the implicit inverse problem actually being targeted, providing insight into its structure and enabling a qualitative comparison with other kernel-based adversarial estimators, as discussed in Section 4.6.

In our derivation of the finite-sample bounds, we make a number of assumptions, which we now introduce and discuss. Our first assumption is a realizability condition:

Assumption 4.1. *There exists a solution to equation (4.2) of the form $I_0 h$ where $h \in \mathcal{H}$;*

equivalently, the set $\mathcal{H}_0 := \{h \in \mathcal{H} : I_0 h \text{ solves equation (4.2)}\}$ is nonempty.

When $I_0 \neq 1$, Assumption 4.1 does not require a solution to equation (4.2) to lie in $\mathcal{H}$ itself. This is important in cases where ρ has a structural zero on $\mathcal{X} \setminus \mathcal{X}_0$: then no solution lies in $\mathcal{H}$ unless the RKHS contains functions that vanish exactly on $\mathcal{X} \setminus \mathcal{X}_0$. Since most RKHSs used in practice do not have this property, existing KRAS estimators from the literature—which implicitly assume $\mathcal{H}$ contains the exact solution—are generally unsuitable in such cases. This motivates our modified KRAS estimators, which explicitly accommodate instances with structural zeros.

When $\mathcal{X}_0 = \mathcal{X}$, so that $I_0 \equiv 1$, the realizability condition above reduces to the standard assumption used in Dikkala et al. (2020) and Ghassami et al. (2022). Note that the condition $\mathcal{H}_0 \cap \mathcal{N}(\mathcal{T}) = \{0\}$ required for the bound on the RMSE computed by Dikkala et al. (2020) and Ghassami et al. (2022) to be non-vacuous rules out that $\mathcal{H}_0$ has more than one element. In contrast, we do not assume that $\mathcal{H}_0$ is a singleton. Instead, we allow equation (4.2) to admit multiple solutions and derive error bounds on $\|\tilde{h}_{KRAS} - h_{0,KRAS}^\dagger\|_2$, where

$$h_{0,KRAS}^\dagger := \arg \min_{h \in \mathcal{H}_0} \|h\|_{\mathcal{H}}^2. \quad (4.13)$$

To ensure that $h_{0,KRAS}^\dagger$ in (4.13) is well-defined—meaning the argmin exists and is unique—we will make the following assumption.

Assumption 4.2. $\mathcal{W}$ is compact and the kernel function $K_{\mathcal{H}} : \mathcal{W} \times \mathcal{W} \rightarrow \mathbb{R}$ is continuous.

In Lemma C.2 of Appendix C.2 we show that under this assumption, the operator $I_0^{op} \circ \mathcal{T}|_{\mathcal{H}} : \mathcal{H} \rightarrow \mathcal{L}_2^2(P_Z)$ is bounded and linear, where $I_0^{op} : \mathcal{L}^2(P_Z) \rightarrow \mathcal{L}^2(P_Z)$ is the operator $(I_0^{op}g)(z) := I_0(x) \cdot g(z)$ where $z = (x, z')$, $\mathcal{T}|_{\mathcal{H}} : \mathcal{H} \rightarrow \mathcal{L}^2(P_Z)$ is the restriction of the conditional mean operator to $\mathcal{H}$, and the boundedness of $I_0^{op} \circ \mathcal{T}|_{\mathcal{H}}$ is with respect to the topology defined by the norm of $\mathcal{H}$. The boundedness and linearity of $I_0^{op} \circ \mathcal{T}|_{\mathcal{H}}$ implies that

its null space is closed and linear, so the projection of any element of $\mathcal{H}_0$ onto it exists. The residual from such projection is the unique element of $\mathcal{H}_0$ with minimum $\mathcal{H}$ -norm.

The constant B in the class $\mathcal{H}_B$ in the definition of $\tilde{h}_{KRAS}$ is any constant satisfying $\|h_{0,KRAS}^\dagger\|_{\mathcal{H}} < B$.

It is important to note that under Assumption 4.2, the integral operator $T_{\mathcal{H}} : \mathcal{L}^2(P_W) \rightarrow \mathcal{L}^2(P_W)$ defined for any $h \in \mathcal{L}^2(P_W)$, as

$$T_{\mathcal{H}}h(w) := \int K_{\mathcal{H}}(w, \tilde{w}) h(\tilde{w}) dP_W(\tilde{w}) \quad (4.14)$$

is a self-adjoint, Hilbert-Schmidt operator (see Theorem 2.34 of Carrasco et al. (2007)), and therefore compact, with respect to the topology defined by the $\mathcal{L}^2(P_W)$ -norm in the domain and the image. On the other hand, it is well known that the range of $T_{\mathcal{H}}$ is included in $\mathcal{H}$. However, the operator $T_{\mathcal{H}} : \mathcal{L}^2(P_W) \rightarrow \mathcal{H}$ is not compact with respect to the topology defined by the $\mathcal{H}$ -norm in the image.

We introduce our two source conditions by first noting that the surrogate regularized problem that the KRAS estimator $\tilde{h}_{KRAS}$ targets, where for simplicity we take $c = 1$, is

$$\begin{aligned} h_{\lambda_{\mathcal{H}},KRAS}^\dagger &:= \arg \min_{h \in \mathcal{H}_B} \left\{ \frac{1}{4} \|\rho - I_0^{op} \circ \mathcal{T}|_{\mathcal{H}} h\|_2^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2 \right\} \\ &= \arg \min_{h \in \mathcal{H}} \left\{ \frac{1}{4} \|\rho - I_0^{op} \circ \mathcal{T}|_{\mathcal{H}} h\|_2^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2 \right\}. \end{aligned} \quad (4.15)$$

For a proof of the second equality see Lemma C.3. Now, in Lemma C.2 of Appendix C.2 we show that the adjoint $(I_0^{op} \circ \mathcal{T}|_{\mathcal{H}})^* : \mathcal{L}^2(P_Z) \rightarrow \mathcal{H}$ of $I_0^{op} \circ \mathcal{T}|_{\mathcal{H}}$ is given by $T_{\mathcal{H}} \circ \mathcal{T}^* \circ I_0^{op}$ where, recall, $\mathcal{T}^* : \mathcal{L}^2(P_Z) \rightarrow \mathcal{L}^2(P_W)$ is the adjoint of the conditional mean operator $\mathcal{T} : \mathcal{L}^2(P_W) \rightarrow \mathcal{L}^2(P_Z)$.

Identifying in Theorem 4.2 the operator $\mathcal{A}$ with $I_0^{op} \circ \mathcal{T}|_{\mathcal{H}}$, the target function h_0 with $h_{0,KRAS}^\dagger$, and the regularized solution h_λ with $h_{\lambda_{\mathcal{H}},KRAS}^\dagger$, and noting that for any $h \in \mathcal{H}$,

$\mathcal{T}^* \circ I_0^{op} \circ I_0^{op} \circ \mathcal{T}|_{\mathcal{H}} h = \mathcal{T}^* \circ I_0^{op} \circ \mathcal{T} h$ naturally leads to the following source condition

Assumption 4.3. *There exists $\beta > 0$ such that $h_{0,KRAS}^\dagger = (T_{\mathcal{H}} \circ \mathcal{T}^* \circ I_0^{op} \circ \mathcal{T})^{\beta/2} h^*$ for some $h^* \in \mathcal{H}$. Equivalently,*

$$h_{0,KRAS}^\dagger = (T_{\mathcal{H}} \circ \mathcal{T}^* \circ I_0^{op} \circ \mathcal{T})^{\beta/2} T_{\mathcal{H}}^{1/2} h^{**} \quad \text{for some } h^{**} \in \mathcal{L}^2(P_W). \quad (4.16)$$

To motivate our second source condition, consider the operator $T_{\mathcal{H}}^{1/2} : \mathcal{L}^2(P_W) \rightarrow \mathcal{H}$. This is a surjective operator that is bounded when $\mathcal{H}$ is endowed with the topology induced by the RKHS inner product. Furthermore, when restricted to the orthogonal complement to the null space of $T_{\mathcal{H}}$, the operator $T_{\mathcal{H}}^{1/2}$ is an isometry. That is, $\|T_{\mathcal{H}}^{1/2} h\|_{\mathcal{H}} = \|h\|_2$ for all $h \in \mathcal{N}(T_{\mathcal{H}})^\perp$ (see Lemma C.1 (d) in Appendix C.2). In addition, the linear operator $\widetilde{\mathcal{T}} := I_0^{op} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2} : \mathcal{L}^2(P_W) \rightarrow \mathcal{L}^2(P_Z)$ is bounded (see Lemma C.2 in Appendix C.2). Then, letting $h'_{\lambda_{\mathcal{H}},KRAS} := \arg \min_{h \in \mathcal{L}^2(P_W)} \left\{ \frac{1}{4} \|\rho - \widetilde{\mathcal{T}} h\|_2^2 + \lambda_{\mathcal{H}} \|h\|_2^2 \right\}$, we have that

$$\begin{aligned} h_{\lambda_{\mathcal{H}},KRAS}^\dagger &:= \arg \min_{h \in \mathcal{H}} \left\{ \frac{1}{4} \|\rho - I_0^{op} \circ \mathcal{T}|_{\mathcal{H}} h\|_2^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2 \right\} \\ &= T_{\mathcal{H}}^{1/2} \left[\arg \min_{h' \in \mathcal{N}(T_{\mathcal{H}})^\perp} \left\{ \frac{1}{4} \left\| \rho - I_0^{op} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2} h' \right\|_2^2 + \lambda_{\mathcal{H}} \left\| T_{\mathcal{H}}^{1/2} h' \right\|_{\mathcal{H}}^2 \right\} \right] \\ &= T_{\mathcal{H}}^{1/2} \left[\arg \min_{h' \in \mathcal{N}(T_{\mathcal{H}})^\perp} \left\{ \frac{1}{4} \left\| \rho - \widetilde{\mathcal{T}} h' \right\|_2^2 + \lambda_{\mathcal{H}} \|h'\|_2^2 \right\} \right] \\ &= T_{\mathcal{H}}^{1/2} \left[\arg \min_{h' \in \mathcal{L}^2(P_W)} \left\{ \frac{1}{4} \left\| \rho - \widetilde{\mathcal{T}} h' \right\|_2^2 + \lambda_{\mathcal{H}} \|h'\|_2^2 \right\} \right] := T_{\mathcal{H}}^{1/2} h'_{\lambda_{\mathcal{H}},KRAS}, \end{aligned}$$

where the second to last equality follows because the minimum over $h' \in \mathcal{L}^2(P_W)$ of the expression in curly brackets is attained at an element of $\mathcal{N}(T_{\mathcal{H}})^\perp$. Note that the first line in the display indicates that the KRAS regularized problem combines a $\mathcal{L}^2$ -loss with an RKHS-norm penalization. The second to last equality establishes that, by applying a suitable

transformation to the underlying operator $\mathcal{T}$, we can rewrite the KRAS problem so that the penalization is also in $\mathcal{L}^2$. This reframing makes the KRAS estimators and the LRAS estimators directly comparable as solutions to regularization problems with the same norm structure but different operators. Specifically, $\tilde{h}_{KRAS}$ is the image—under the isometry $T_{\mathcal{H}}^{1/2}$ —of the Tikhonov-regularized solution $h'_{\lambda_{\mathcal{H}}, KRAS}$ (with an $\mathcal{L}^2$ penalty) to the operator equation

$$\tilde{\mathcal{T}}h = \rho. \quad (4.17)$$

This observation has two important implications. First, since $T_{\mathcal{H}}^{1/2}$ is a compact operator (see Lemma C.1 (f) in Appendix C.2), regularization in the RKHS norm—as used in the computation of the KRAS estimators—amounts to solving a more severely ill-posed problem than the original integral equation (4.2), because the modified operator $\tilde{\mathcal{T}}$ smooths more aggressively than $\mathcal{T}$. Second, by assuming a β -source condition on the minimal $\mathcal{L}^2(P_W)$ -norm solution $h'_{0, KRAS}$ to equation (4.17), we can derive bounds on the regularization bias $\|h'_{\lambda_{\mathcal{H}}, KRAS} - h'_{0, KRAS}\|_2$. Lemma C.4 in Appendix C.2 establishes that $h_{0, KRAS}^\dagger = T_{\mathcal{H}}^{1/2} h'_{0, KRAS}$. This then implies that with such source condition we can also control the regularization bias $\|h_{\lambda_{\mathcal{H}}, KRAS}^\dagger - h_{0, KRAS}^\dagger\|_2$ of the regularized solution to the original problem $\mathcal{T}I_0h = \rho(Z)$.

The first point is key to understanding the distinction between RMSE bounds for LRAS, i.e. RAS estimators over RKHS classes $\mathcal{H}$ and $\mathcal{G}$ that use $\mathcal{L}^2(\mathbb{P}_n)$ penalization, versus KRAS estimators, i.e. those that employ RKHS penalization—a topic we return to in Section 4.6.1. The second point, in turn, motivates our announced second source condition. To state it, we first note that the adjoint $\tilde{\mathcal{T}}^* : \mathcal{L}^2(P_Z) \rightarrow \mathcal{L}^2(P_W)$ of $\tilde{\mathcal{T}}$ satisfies $\tilde{\mathcal{T}}^* = T_{\mathcal{H}}^{1/2} \circ \mathcal{T}^* \circ I_0^{op}$.

Assumption 4.4. *There exists $\beta > 0$ and $h^{**} \in \mathcal{L}^2(P_W)$ such that*

$$h_{0, KRAS}^\dagger = T_{\mathcal{H}}^{1/2} \circ \left(\tilde{\mathcal{T}}^* \circ \tilde{\mathcal{T}} \right)^{\beta/2} h^{**} = T_{\mathcal{H}}^{1/2} \circ \left(T_{\mathcal{H}}^{1/2} \circ \mathcal{T}^* \circ I_0^{op} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2} \right)^{\beta/2} h^{**}. \quad (4.18)$$

Assumption 4.4 is tantamount to requiring the standard β -source condition on $h'_{0,KRAS}$ with respect to the operator $\widetilde{\mathcal{T}}$. When $\widetilde{\mathcal{T}}$ is compact—a property that is satisfied when W and Z do not share common components—this requires that $h'_{0,KRAS}$ decays rapidly in the directions corresponding to small singular values of $\widetilde{\mathcal{T}}$, with larger values of β implying greater smoothness and better alignment with the stable directions of $\widetilde{\mathcal{T}}$. The final application of $T_{\mathcal{H}}^{1/2}$ in Assumption 4.4 further smooths $h'_{0,KRAS}$ with respect to the singular value decomposition of the integral operator $T_{\mathcal{H}}$. Thus, Assumption 4.4 encodes smoothness of $h_{0,KRAS}^{\dagger}$ relative to both the operator $\widetilde{\mathcal{T}}$ and the geometry of the RKHS $\mathcal{H}$.

As we will see in the next theorem, Assumptions 4.3 and 4.4 lead to identical bounds on the regularization bias. The assumptions coincide when $\beta = 2$ but impose different smoothness requirements otherwise. We find Assumption 4.4 particularly useful for comparisons with RAS estimators that penalize using the $\mathcal{L}^2(\mathbb{P}_n)$ -norm, as it posits a β -source condition related to the operator $\widetilde{\mathcal{T}}$ acting between $\mathcal{L}^2$ spaces, rather than for an operator whose domain is the RKHS $\mathcal{H}$.

Theorem 4.3 (Regularization Bias). *Suppose that Assumptions 4.1 and 4.2 hold. If, in addition, Assumption 4.3 or Assumption 4.4 hold, then there exist constants c_1^* and c_2^* such that*

$$\|h_{\lambda_{\mathcal{H}},KRAS}^{\dagger} - h_{0,KRAS}^{\dagger}\|_{\mathcal{H}}^2 \leq c_1^* \lambda_{\mathcal{H}}^{\min\{\beta,2\}} \text{ and } \|\mathcal{T}I_0(h_{\lambda_{\mathcal{H}},KRAS}^{\dagger} - h_{0,KRAS}^{\dagger})\|_2^2 \leq c_2^* \lambda_{\mathcal{H}}^{\min\{\beta+1,2\}}.$$

Under Assumption 4.3, Theorem 4.3 follows by a direct application of Theorem 4.2. In Appendix C.1, we provide the proof of Theorem 4.3 under Assumption 4.4. Having established conditions under which the regularization bias can be controlled, we now turn to the

additional assumptions required to bound the estimation error. Our derivation follows the framework of [Bennett et al. \(2023b\)](#), and therefore relies on conditions closely aligned with theirs. The resulting bounds will be expressed in terms of the critical radius of appropriately defined subsets derived from $\mathcal{G}$. In the course of our derivations, we identified a few instances—particularly in the application of a Lemma in [Foster and Syrgkanis \(2023\)](#) (see Appendix C.2)—where the assumptions made by [Bennett et al. \(2023b\)](#) on these subsets can be slightly relaxed. The first such assumption, stated below, slightly weakens a similar condition made by all existing analyses of RAS estimators as we discuss subsequently:

Assumption 4.5 (Local closeness and boundedness of $\mathcal{T}$). *There exists $B > \|h_{0,KRAS}^\dagger\|_{\mathcal{H}}$ and $B' > 0$ such that if $\|h\|_{\mathcal{H}} \leq 2B$ then $\mathcal{T}h \in \mathcal{G}$ and $\|\mathcal{T}h\|_{\mathcal{G}} \leq B'$.*

The local closeness condition, i.e., that $\mathcal{T}h \in \mathcal{G}$ for h in some sufficiently large RKHS ball is used in the derivation of the error bounds to connect $h_{\lambda_{\mathcal{H}},KRAS}^\dagger$ with the population version of the regularized optimization problem targeted by the KRAS estimation procedure. Specifically, by part 1 of Proposition 4.2, it follows that under this assumption,

$$h_{\lambda_{\mathcal{H}},KRAS}^\dagger = \arg \min_{h \in \mathcal{H}} \max_{g \in \mathcal{G}} \mathbb{E} \{ m(O; g) - I_0(X)r(O)h(W)g(Z) - c^2 \cdot I_0(X)g(Z)^2 \} + \lambda_{\mathcal{H}}\|h\|_{\mathcal{H}}. \quad (4.19)$$

Note that the right-hand side of (4.19) is a population optimization problem that parallels the empirical optimization problem solved by the proposed KRAS estimator $\tilde{h}_{KRAS}$, except that the estimation also includes the penalty term $\lambda_{\mathcal{G}}\|g\|_{\mathcal{G}}^2$ in the inner maximization part.

The boundedness condition in Assumption 4.5, i.e. that if $\|h\|_{\mathcal{H}} \leq 2B$ then $\mathcal{T}h \in \mathcal{G}$ and $\|\mathcal{T}h\|_{\mathcal{G}} \leq B'$, is invoked to ensure that $\|\mathcal{T}(\tilde{h}_{KRAS} - h_{0,KRAS}^\dagger)\|_{\mathcal{G}} = O(1)$ which, in turn, is used to control the influence of the penalization $\lambda_{\mathcal{G}}\|g\|_{\mathcal{G}}^2$ of the inner maximization in the computation of (4.12). In their analysis of the KRAS estimator, [Dikkala et al. \(2020\)](#) and [Ghassami et al. \(2022\)](#) imposed the related global closeness that $\mathcal{T}(h - h_{0,KRAS}^\dagger) \in \mathcal{G}$ for all

$h \in \mathcal{H}$ —together with the related Lipchitz/boundedness condition that $\|\mathcal{T}(h - h_{0,KRAS}^\dagger)\|_{\mathcal{G}} \leq C\|h - h_{0,KRAS}^\dagger\|_{\mathcal{H}}$ for some constant C . [Bennett et al. \(2023b\)](#) also adopted a similar condition when deriving error bounds for the LRAS estimator. When, as we assume throughout, $h_{0,KRAS}^\dagger \in \mathcal{H}$ this requirement is stronger than Assumption 4.5. [Dikkala et al. \(2020\)](#) discuss conditions under which their assumption, and consequently our Assumption 4.5, holds. Nonetheless, Assumption 4.5 remains, in our view, the most conceptually opaque and difficult to justify among those used in the analysis of RAS estimators.

We will also make the following assumption:

Assumption 4.6. $\mathcal{Z}$ is compact and the kernel function $K_{\mathcal{G}} : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ is continuous

Assumption 4.2 implies that $k_1 := \{\sup_{w \in \mathcal{W}} K_{\mathcal{H}}(w, w)\}^{1/2} < \infty$ and Assumption 4.6 that $k_2 := \{\sup_{z \in \mathcal{Z}} K_{\mathcal{G}}(z, z)\}^{1/2} < \infty$. An important implication that is used repeatedly in the derivation of the error bounds, is that, if we let $\mathcal{F}_B := \{f \in \mathcal{F} : \|f\|_{\mathcal{F}} \leq B\}$ where $\mathcal{F}$ is either $\mathcal{H}$ or $\mathcal{G}$ and $\|f\|_{\mathcal{F}}$ is the norm in the corresponding RKHS, then since $\|f\|_{\infty} \leq \|f\|_{\mathcal{F}} \{\sup_{v \in \mathcal{V}} K_{\mathcal{F}}(v, v)\}^{1/2}$, the class $\mathcal{F}_B$ is *uniformly bounded* in the sense that $\|f\|_{\infty} \leq C$ for some universal constant C and for all $f \in \mathcal{F}_B$.

To state the results that provide the announced finite-sample bounds, we introduce additional notation and review important concepts from the theory of concentration inequalities. Let V denote a random vector and let $\mathcal{F}$ denote a class of real-valued measurable functions $f : \mathcal{V} \rightarrow \mathbb{R}$ included on a normed space. The class $\mathcal{F}$ is called *star-shaped* if for any $f \in \mathcal{F}$ and $\alpha \in [0, 1]$, the scaled function αf also belongs to $\mathcal{F}$. We define $\text{star}(\mathcal{F}) := \{\alpha f : f \in \mathcal{F}, \alpha \in [0, 1]\}$, which is the smallest star-shaped class that contains all the functions in $\mathcal{F}$. For any function $f^* \in \mathcal{F}$, we define $\mathcal{F} - f^* := \{f - f^* : f \in \mathcal{F}\}$.

For a given $\delta > 0$, the *localized Rademacher complexity of radius δ* of a function class $\mathcal{F}$

is defined as

$$\mathcal{R}_n(\mathcal{F}, \delta) := \mathbb{E}_{\varepsilon, V_1, \dots, V_n} \left[\sup_{f \in \mathcal{F}: \|f\|_2 \leq \delta} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i f(V_i) \right| \right],$$

where $V_1, \dots, V_n$ are independent and identically distributed copies of a random variable V and $\varepsilon_1, \dots, \varepsilon_n$ are independent Rademacher random variables (Bartlett et al., 2005), independent of $V_1, \dots, V_n$. The localized Rademacher complexity of the function class $\mathcal{F} - f^*$ is a measure of the complexity of a class of functions within a neighborhood of a given target function f^* . This measure focuses on functions that are *close* to the target, thereby avoiding unnecessary increases in complexity by ignoring functions that are far from the target. Intuitively, this measure captures the capacity of the functions in $\mathcal{F}$ in a neighborhood of f^* to fit noise sequences of size n .

The *critical radius* of a function class $\mathcal{F}$, uniformly bounded by b , is defined as the smallest positive radius δ for which the following inequality holds $\mathcal{R}_n(\mathcal{F}, \delta) \leq \frac{\delta^2}{b}$. We are now ready to state our first main result.

Theorem 4.4. *Suppose $m(o; g) = I_0(x)s(o)g(z)$ for some fixed real-valued function s satisfying $\|s\|_\infty < \infty$. Suppose Assumptions 4.1, 4.2, 4.5, 4.6 hold and either Assumption 4.3 or 4.4 hold when $P = P_0$. Suppose that the constant B in Assumption 4.5 is taken large enough so that $\|h_{0,KRAS}^\dagger\|_{\mathcal{H}} \leq B$. Let $B_1 = B'/c^2$. Define the function class $I_0 \cdot \mathcal{G}_{B_1} := \{I_0 \cdot g : g \in \mathcal{G}_{B_1}\}$ and let δ_n be an upper bound on the critical radius of $I_0 \cdot \mathcal{G}_{B_1}$. If $\lambda_{\mathcal{H}} \leq 1$, and $\lambda_{\mathcal{G}} \geq 157c^2\delta_n^2/B_1^2$, then with probability at least $\kappa(\delta_n) := 1 - [36 + 18 \cdot I\{\delta_n < 5b/6\} \cdot \lfloor \log(6/5) \rfloor^{-1} \log(b/\delta_n)] \exp(-c_1 n \delta_n^2)$ where $c_1 = \{8b^2(1 + 17e)\}^{-1}$ and $b = k_2 B_1$, it holds that $\hat{h}_{KRAS} := I_0 \tilde{h}_{KRAS}$ and $h_{0,KRAS} := I_0 h_{0,KRAS}^\dagger$ satisfy*

$$\left\| \hat{h}_{KRAS} - h_{0,KRAS} \right\|_2^2 \leq c_2 \left(\frac{\delta_n^2}{\lambda_{\mathcal{H}}} + \frac{\lambda_{\mathcal{G}}}{\lambda_{\mathcal{H}}} + \lambda_{\mathcal{H}}^{\min\{\beta, 1\}} \right), \quad (4.20)$$

and

$$\left\| \mathcal{T}(\hat{h}_{KRAS} - h_{0,KRAS}) \right\|_2^2 \leq c_3 \left(\delta_n^2 + \lambda_{\mathcal{G}} + \lambda_{\mathcal{H}}^{\min\{\beta+1, 2\}} \right), \quad (4.21)$$

where c_2 and c_3 are universal constants that depend neither on n nor on δ_n .

Notice that the bounds in (4.20) and (4.21) depend on the critical radius of only one function class, namely the class $I_0 \cdot \mathcal{G}_{B_1}$. When $I_0 = 1$, this reduces to the critical radius of $\mathcal{G}_{B_1}$ itself, which is well understood for many function classes, including various RKHSs; see, for example, Chapters 13 and 14 of [Wainwright \(2019\)](#).

Theorem 4.4 provides bounds on the weak error $\|\mathcal{T}(\hat{h}_{KRAS} - h_{0,KRAS})\|_2$ as well as the RMSE, $\|\hat{h}_{KRAS} - h_{0,KRAS}\|_2$, for the special case in which $m(o; g) = I_0(x)s(o)g(z)$. Our next result gives bounds on the weak error and the RMSE when the map $g \mapsto m(o; g)$ is an arbitrary linear map for each fixed $o \in \mathcal{O}$ satisfying the following additional condition:

Assumption 4.7.

1. There exists $B_2 > 0$ such that, for any $g \in \mathcal{L}^2(P_Z)$, it holds that $\|m(\cdot; g)\|_2 \leq B_2 \|I_0 g\|_2$.
2. The function class $m \circ \mathcal{G} := \{o \mapsto m(o; g) : g \in \mathcal{G}\}$ can be equipped with a metric d under which it is separable. Moreover, d is controlled by the RKHS norm in the sense that $d(m(\cdot; g), m(\cdot; g')) \leq \|g - g'\|_{\mathcal{G}}$ for all $g, g' \in \mathcal{G}$ and it dominates the metric induced by the sup-norm on $\mathcal{O}$, meaning there exists a constant $B_3 > 0$ such that $\sup_{o \in \mathcal{O}} |f(o) - f'(o)| \leq B_3 \cdot d(f, f')$ for all $f, f' \in m \circ \mathcal{G}$.

In Appendix C.1, we verify that Assumption 4.7 holds in Example 4.2 under suitable conditions of the map q and the data-generating process.

Theorem 4.5. Suppose Assumptions 4.1, 4.2, 4.5, 4.6, 4.7 hold and either Assumption 4.3 or 4.4 hold for $P = P_0$. Suppose that the constant B in Assumption 4.5 is taken large enough

so that $\|h_{0,KRAS}^\dagger\|_{\mathcal{H}} \leq B$. Let $B_1 = B'/c^2$. Define $m \circ \mathcal{G}_{B_1} := \{o \mapsto m(o; g) : g \in \mathcal{G}_{B_1}\}$, and let the function class $I_0 \cdot \mathcal{G}_{B_1}$ be defined as in Theorem 4.4. Let δ_n be an upper bound on the critical radius of $I_0 \cdot \mathcal{G}_{B_1}$ and of $\text{star}(m \circ \mathcal{G}_{B_1})$. If $\lambda_{\mathcal{H}} \leq 1$, and $\lambda_{\mathcal{G}} \geq 157c^2\delta_n^2/B_1^2$, then inequalities (4.20) and (4.21) hold for some universal constants c_2 and c_3 , with probability at least as $\kappa(\delta_n)$ as given in Theorem 4.4 with c_1 as defined in that Theorem, but with $b := B_1 \max\{k_2, B_3\}$, where B_3 is the constant specified in Assumption 4.7.

The distinction between Theorems 4.4 and 4.5 is that in Theorem 4.4, δ_n appearing in the bounds (4.20) and (4.21) is an upper bound on the critical radius of the class $I_0 \cdot \mathcal{G}_{B_1}$, whereas in Theorem 4.5, δ_n stands for an upper bound of the critical radius of $I_0 \cdot \mathcal{G}_{B_1}$ and $\text{star}(m \circ \mathcal{G}_{B_1})$. In the special case $m(o; g) = g\{q(z)\}$ for a measurable map $q : \mathcal{Z} \rightarrow \mathcal{Z}$, the following result shows that the critical radius of $\text{star}(m \circ \mathcal{G}_{B_1})$ is dominated by that of $\mathcal{G}_{B_1}$.

Lemma 4.1. *Suppose $m(o; g) = g\{q(z)\}$ for some measurable map $q : \mathcal{Z} \rightarrow \mathcal{Z}$, where z is a subvector of o . Let δ_n and δ'_n denote the critical radii of $\text{star}(m \circ \mathcal{G}_{B_1})$ and $\mathcal{G}_{B_1}$, respectively. Then, there exists a universal constant c_1'' such that $\delta_n \leq c_1'' \delta'_n$.*

Remark 4.1. *When $\delta_n \geq 5b/6$ for $b = k_2 B_1$ ($= B_1 \max\{k_2, B_3\}$), the bounds on $\|\hat{h}_{KRAS} - h_{0,KRAS}\|_2^2$ and $\|\mathcal{T}(\hat{h}_{KRAS} - h_{0,KRAS})\|_2^2$ provided in Theorem 4.4 (Theorem 4.5) hold with probability at least $1 - 36 \exp(-c_1 n \delta_n^2)$. On the other hand, if $c_0 \log(\log(n))/n \leq \delta_n^2 < (5b/6)^2$ with $c_0 = \max\{b^2/\log(\log(3)), 2/c_1\}$, and $n \geq 3$, the bounds in Theorem 4.4 and Theorem 4.5 remain valid with probability at least $1 - 18 \cdot \{\log(6/5)\}^{-1} \exp(-\frac{c_1}{2} n \delta_n^2)$, (see Appendix C.1), which tends to 1 as $n \rightarrow \infty$ because $n \delta_n^2 \geq c_0 \log(\log(n)) \rightarrow \infty$. The condition $c_0 \log(\log(n))/n \leq \delta_n^2$ holds for a sufficiently large n , in particular, when $\mathcal{G}$ is an RKHS with Gaussian kernel because in such case $\delta_n^2 = O(\log(n)/n)$.*

Theorems 4.4 and 4.5 have the following important corollary.

Corollary 4.2. *Suppose that the conditions of Theorems 4.4 or 4.5 hold. If the penalization parameters are chosen as $\lambda_{\mathcal{H}} = O(\delta_n^{2/(1+\min\{\beta,1\})})$ and $\lambda_{\mathcal{G}} = O(\delta_n^2)$, then*

$$\left\| \hat{h}_{KRAS} - h_{0,KRAS} \right\|_2 = O_p \left(\delta_n^{\frac{\min\{\beta,1\}}{1+\min\{\beta,1\}}} \right) \quad \text{and} \quad \left\| \mathcal{T}(\hat{h}_{KRAS} - h_{0,KRAS}) \right\|_2 = O_p(\delta_n).$$

Recall from Theorem 4.1 that in order for the debiased machine learning estimator $\hat{\theta}$ of $\theta(P_0)$ to be asymptotically linear, the remainder term R_n must be of order $o_p(n^{-1/2})$. In view of Corollary 4.2, it follows that if: (i) $\tilde{h}$ and $\tilde{g}$ are the KRAS estimators $\tilde{h}_{KRAS}$ and $\tilde{g}_{KRAS}$, (ii) the conditions of Theorems 4.4 or 4.5 hold for the problem of estimating the solutions $h_{0,KRAS} := I_0 h_{0,KRAS}^\dagger$ and $g_{0,KRAS} := I_0 g_{0,KRAS}^\dagger$ of equations (4.2) and (4.3), (iii) $h_{0,KRAS}^\dagger$ and $g_{0,KRAS}^\dagger$ satisfy Assumption 4.3 or Assumption 4.4 (for their corresponding operators) with parameters β_h^{KRAS} and β_g^{KRAS} , (iv) the relevant critical radius in Theorems 4.4 or 4.5 for the classes used in the estimation of $h_{0,KRAS}$ and $g_{0,KRAS}$ are the same (up to multiplicative constants) and equal to universal constants times δ_n , and (v) the tuning parameters are selected as indicated in Corollary 4.2, then

$$R_n = O(\min\{\|\hat{h}_{KRAS} - h_{0,KRAS}\|_2 \|\mathcal{T}^*(\hat{g}_{KRAS} - g_{0,KRAS})\|_2, \|\hat{g}_{KRAS} - g_{0,KRAS}\|_2 \|\mathcal{T}(\hat{h}_{KRAS} - h_{0,KRAS})\|_2\}) = O_p \left(\delta_n^{1 + \frac{\beta^{KRAS}}{1 + \beta^{KRAS}}} \right),$$

where $\beta^{KRAS} = \min\{\max\{\beta_h^{KRAS}, \beta_g^{KRAS}\}, 1\}$. In particular, if the RKHSs used in the internal maximization in both estimation problems are Gaussian, then $\delta_n = O(\sqrt{\log(n)/n})$, which yields $R_n = o_p(n^{-1/2})$. In contrast, if they are d -dimensional RKHSs with Sobolev kernels of order $\nu = (l+d)/2$ where l is a positive integer, then $\delta_n = O(n^{-\nu/(2\nu+d)})$, which then yields $R_n = o_p(n^{-1/2})$ whenever $d < l$ and $\beta \in (d/l, 1]$. Thus, we observe double robustness

in the smoothness of the minimum RKHS–norm targets $h_{0,KRAS}^\dagger$ and $g_{0,KRAS}^\dagger$. Specifically, even if one of $h_{0,KRAS}^\dagger$ or $g_{0,KRAS}^\dagger$ has low smoothness (i.e., the corresponding parameter of the source condition in Assumption 4.3 or Assumption 4.4 is small), the remainder term R_n can still be controlled as long as the other nuisance is sufficiently smooth.

As in the work of Bennett et al. (2023b), we could in principle improve upon the KRAS estimators $\tilde{h}_{KRAS}$ and $\tilde{g}_{KRAS}$ by using an iterated Tikhonov strategy. The details of such procedure will be reported elsewhere. We conclude with a final comment on the computational implementation of our proposed estimator to ensure it satisfies the conditions required by Theorems 4.4 and 4.5.

Remark 4.2. *The errors bounds established in Theorems 4.4 and 4.5 require that the estimator $\tilde{h}_{KRAS}$ has RKHS norm at most B , where B is the constant specified in Assumption 4.5. While this constant is not typically known in practice, one can select a sufficiently large B and solve the optimization problem in (4.12) over the constrained class $\mathcal{H}_B$. The computation of the solution in the restricted class $\mathcal{H}_B$ reduces to the problem of optimizing a quadratic objective function in $\mathbb{R}_n$ subject to a quadratic constraint. As such, enforcement of the constraint $\|\tilde{h}_{KRAS}\|_{\mathcal{H}} < B$ can be conducted as follows. First a solution to the unconstrained optimization problem is found. If the solution satisfies the constraint, then this is the solution to the constrained problem. If not, the unconstrained optimization problem is altered as indicated in Exercise 4.22 of Boyd (2004), and the unconstrained solution to this new problem agrees with the solution to the original constrained problem.*

4.6 Contrasting Kernel based Adversarial Estimation Strategies

In this section we investigate the advantages and limitations of the three prominent recently proposed kernel-based adversarial approaches: (i) LRAS estimators, i.e. kernel adversarial stabilized estimators with $\mathcal{L}^2(\mathbb{P}_n)$ penalization; (ii) KRAS estimators, i.e. kernel adver-

serial stabilized estimators with RKHS penalization; and (iii) the kernel maximal moment restriction (KMMR) estimators studied in [Muandet et al. \(2020a\)](#); [Mastouri et al. \(2021\)](#); [Kallus et al. \(2021\)](#); and [Zhang et al. \(2023\)](#). We compare these methods in terms of the assumptions they impose (see also Appendix C.3.2), the scope and sharpness of their available finite-sample error bounds, and—most importantly—the nature of the inverse problems they implicitly solve. This perspective brings into focus the trade-offs between statistical convergence rates, computational stability, and the degree of ill-posedness each method must confront. To keep the notation light, we assume throughout this section that $I_0 = 1$.

4.6.1 $\mathcal{L}^2(\mathbb{P}_n)$ -penalized Kernel Regularized Adversarial Stabilized Estimators

An LRAS estimator uses RKHS classes $\mathcal{H}$ and $\mathcal{G}$, and $\mathcal{L}^2(\mathbb{P}_n)$ -penalization:

$$\begin{aligned} \tilde{h}_{LRAS} := \arg \min_{h \in \mathcal{H}_D} \max_{g \in \mathcal{G}} & \left[\mathbb{E}_n \left\{ m(O; g) - r(O)h(W)g(Z) - r(X)g(Z)^2 \right\} \right. \\ & \left. - \lambda_{\mathcal{G}} \mathbb{E}_n \left\{ g(Z)^2 \right\} + \lambda_{\mathcal{H}} \mathbb{E}_n \left\{ h(W)^2 \right\} \right], \end{aligned}$$

for some constant D and where, for simplicity, we take $c = 1$. When $\lambda_{\mathcal{G}}$ is negligible in relation to $\lambda_{\mathcal{H}}$, and the closeness condition (4.8) holds, the LRAS estimator aims at the solution to the surrogate regularized problem:

$$h_{\lambda_{\mathcal{H}}, LRAS} := \arg \min_{h \in \mathcal{H}_D} \|\mathcal{T}h - \rho\|_2^2 + \lambda_{\mathcal{H}} \|h\|_2^2.$$

[Bennett et al. \(2023b\)](#) derived bounds for $\|\tilde{h}_{LRAS} - h_{0, LRAS}\|_2^2$ and $\|\mathcal{T}(\tilde{h}_{LRAS} - h_{0, LRAS})\|_2^2$ for estimators with $\lambda_{\mathcal{G}} = 0$ where $h_{0, LRAS} := \arg \min_{h \in \mathcal{L}^2(P_W): h \text{ solves (4.2)}} \|h\|_2^2$ under the critical

assumption:

$$\arg \min_{h \in \mathcal{L}^2(P_W)} \|\mathcal{T}h - \rho\|_2^2 + \lambda_{\mathcal{H}} \|h\|_2^2 \in \mathcal{H}_D \quad \text{for all } \lambda_{\mathcal{H}} \text{ in a neighborhood of 0.} \quad (4.22)$$

Under this assumption, $h_{\lambda_{\mathcal{H}}, LRAS} = \arg \min_{h \in \mathcal{L}^2(P_W)} \|\mathcal{T}h - \rho\|_2^2 + \lambda_{\mathcal{H}} \|h\|_2^2$. To invoke Theorem 4.2 to derive bounds for the regularization errors $\|h_{\lambda_{\mathcal{H}}, LRAS} - h_{0, LRAS}\|_2$ and $\|\mathcal{T}(h_{\lambda_{\mathcal{H}}, LRAS} - h_{0, LRAS})\|_2$ these authors imposed the β -source condition that posits that there exists $\beta > 0$ such that

$$h_{0, LRAS} = (\mathcal{T}^* \mathcal{T})^{\beta/2} h^{**} \quad \text{for some } h^{**} \in \mathcal{L}^2(P_W). \quad (4.23)$$

Unlike for KRAS estimators, establishing bounds for LRAS does not require the realizability condition in Assumption 4.1 that posits the existence of a solution to equation (4.2) within $\mathcal{H}$. However, bounds for LRAS estimators rely critically on the validity of condition (4.22)—a condition that the analysis of KRAS estimators does not impose. Thus, the analysis of KRAS estimators requires a more natural assumption that directly places the smoothness requirement—i.e., membership in the RKHS $\mathcal{H}$ —only on the true solution of the integral equation itself. In contrast, the analysis of LRAS estimators demands this smoothness property to hold for the solutions of infinitely many surrogate regularized problems, making it a less natural and generally more stringent condition.

The β -source conditions (4.16) or (4.18) assumed for KRAS estimators also differ from the source condition (4.23) invoked for analyzing LRAS estimators in two key ways. Consider, for instance, the source condition (4.18) made in analyzing KRAS estimators. First, the smoothness assumptions in the conditions (4.18) and (4.23) are placed on potentially different solutions to equation (4.2): in the KRAS case, it is the minimum-norm solution in $\mathcal{H}$, whereas in the LRAS case, it is the minimum-norm solution in $\mathcal{L}^2(P_W)$. Second, the fractional power is taken with respect to different operators: $\widetilde{\mathcal{T}}$ in the KRAS case, and $\mathcal{T}$ in the LRAS case.

For the same value of β , condition (4.18) imposes a stronger smoothness requirement than condition (4.23), due to the presence of the additional operator $T_{\mathcal{H}}^{1/2}$ in the definition of $\widetilde{\mathcal{T}}$. This is not surprising, since—as discussed in Section 4.5—regularization in the RKHS norm is ultimately equivalent to solving the more severely ill-posed problem (4.17) rather than the original problem (4.2).

We note that the finite-sample bounds for estimators that use optimal tuning parameters in Bennett et al. (2023b) have the same functional form as those in Corollary 4.2. However, this similarity should not be taken to imply that the bounds for the KRAS and LRAS estimators achieve the same convergence rate under comparable smoothness levels of their respective minimum-norm solutions. This is so because, qualitatively, a given value of β in Assumption (4.23) corresponds to a smaller β in Assumption (4.18) when representing an equivalent level of smoothness in the solution.

A few additional remarks are in place.

- All existing RAS estimators in the literature are for the case $I_0 = 1$. Our paper appears to be the first that incorporates structural zeros into RAS estimation.
- Existing analyses of LRAS estimators, and more generally, of any RAS estimator regardless of the specific function classes $\mathcal{H}$ and $\mathcal{G}$ used, rely on a closeness and boundedness condition of the type in Assumption 4.5. The assumption is invoked to justify the equality $M_P(h) = \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2$. An interesting open question raised by Proposition 4.2 is whether part 2 of this proposition can be used to relax this assumption.
- Bennett et al. (2023b) also analyzed the iterative Tikhonov version of LRAS estimators. For LRAS with t Tikhonov iterations, they established weak and strong error bounds analogous to those in Theorem 4.4, but with the exponents $\min\{\beta, 2\}$ and $\min\{\beta + 1, 2\}$ replaced by $\min\{\beta, 2t\}$ and $\min\{\beta + 1, 2t\}$, respectively. This is as expected since it is well known that iterative Tikhonov regularization with t iterations offers a tangible

benefit in regularization bias when the smoothness parameter β is between 2 and $2t$. We conjecture that a parallel theoretical analysis of t -iterated Tikhonov KRAS estimators should be doable under either Assumption 4.3 or 4.4. As in the non-iterated case, we anticipate that such an analysis will produce finite-sample bounds with the same structural form as those obtained by Bennett et al. (2023b) for their t -iterated procedure.

- Bennett et al. (2023b) also analyzed an adaptive version of the iterated Tikhonov LRAS estimator, deriving weak and strong error bounds with exponents $\beta + 1$ and β , respectively—thereby fully eliminating the saturation phenomenon and enabling convergence rates to improve without bound as smoothness increases. We conjecture that an analogous result holds for adaptive iterated KRAS estimators. We note that while adaptive procedures represent a notable theoretical advance, their practical implementation introduces additional computational challenges that have yet to be fully examined.

4.6.2 Kernel Adversarial Maximal Moment Restriction Estimators

Like regularized adversarial stabilized estimation strategies, maximal moment restriction methods also leverage the fact that the residual (4.6) has mean zero for any g in a given function class $g \in \mathcal{G} \subset \mathcal{L}^2(P_Z)$. A kernel maximal moment estimator (KMMR) estimator (Muandet et al. (2020a); Mastouri et al. (2021)) is defined as

$$\hat{h}_{KMMR} := \arg \min_{h \in \mathcal{H}} \max_{g \in \mathcal{G}: \|g\|_{\mathcal{G}} \leq 1} [\mathbb{E}_n \{m(O; g) - h(W)g(Z)\}]^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2.$$

where $\mathcal{H}$ and $\mathcal{G}$ are RKHSs with associated norms $\|h\|_{\mathcal{H}}$ and $\|g\|_{\mathcal{G}}$ and $\lambda_{\mathcal{H}} > 0$. Park et al. (2024) adopts a similar strategy to construct solutions to integral equations needed

for estimation of target parameters that fall outside the class defined in (4.1). Like KRAS estimators, KMMR estimators admit a closed-form expression and avoid inversion of ill-posed matrices.

To facilitate the comparison with KRAS estimators we will re-express the surrogate problem that the KMMR estimator targets in terms of the weak error relative to some operator. To do so, we start by noting that an immediate generalization of the results in Zhang et al. (2023) yields:

$$\max_{g \in \mathcal{G}, \|g\|_{\mathcal{G}} \leq 1} [\mathbb{E}\{m(O; g) - h(W)g(Z)\}]^2 = \|T_{\mathcal{G}}(\rho - \mathcal{T}h)\|_{\mathcal{G}}^2$$

where $T_{\mathcal{G}} : \mathcal{L}^2(P_Z) \rightarrow \mathcal{G}$ is the integral operator associated with the kernel $K_{\mathcal{G}}$ that defines the RKHS $\mathcal{G}$, analogously to the operator defined in (4.14). Writing $T_{\mathcal{G}} = T_{\mathcal{G}}^{1/2} \circ T_{\mathcal{G}}^{1/2}$ and noticing that, under Assumption 4.6, $T_{\mathcal{G}}^{1/2} : \mathcal{L}^2(P_Z) \rightarrow \mathcal{G}$ is an isometry when restricted to the orthogonal complement to the null space of $T_{\mathcal{G}}$, and therefore when restricted to the range of $T_{\mathcal{G}}^{1/2}$, we have that $\|T_{\mathcal{G}}(\rho - \mathcal{T}h)\|_{\mathcal{G}}^2 = \|T_{\mathcal{G}}^{1/2}(\rho - \mathcal{T}h)\|_2^2$. Then, recalling that $T_{\mathcal{H}}^{1/2} : \mathcal{L}^2(P_W) \rightarrow \mathcal{H}$ is surjective, we conclude that the regularized solution targeted by the KMMR estimator $\hat{h}_{KMMR}$ is

$$\begin{aligned} h_{\lambda_{\mathcal{H}}, KMMR}^{\dagger} &:= \arg \min_{h \in \mathcal{H}} \max_{g \in \mathcal{G}: \|g\|_{\mathcal{G}} \leq 1} [\mathbb{E}\{m(O; g) - h(W)g(Z)\}]^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2 \\ &= \arg \min_{h \in \mathcal{H}} \|T_{\mathcal{G}}(\rho - \mathcal{T}h)\|_{\mathcal{G}}^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2 \end{aligned} \quad (4.24)$$

$$\begin{aligned} &= T_{\mathcal{H}}^{1/2} \left[\arg \min_{h' \in \mathcal{N}(T_{\mathcal{H}})^{\perp}} \|T_{\mathcal{G}}\rho - T_{\mathcal{G}} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2} h'\|_{\mathcal{G}}^2 + \lambda_{\mathcal{H}} \|h'\|_2^2 \right] \\ &= T_{\mathcal{H}}^{1/2} \left[\arg \min_{h' \in \mathcal{L}^2(P_W)} \|T_{\mathcal{G}}^{1/2}\rho - T_{\mathcal{G}}^{1/2} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2} h'\|_2^2 + \lambda_{\mathcal{H}} \|h'\|_2^2 \right] \\ &:= T_{\mathcal{H}}^{1/2} h'_{\lambda_{\mathcal{H}}, KMMR}. \end{aligned} \quad (4.25)$$

Just as it was the case for the analysis of KRAS estimators, the preceding display implies

two ways of analyzing the integral equation that the KMMR estimators are implicitly targeting. Specifically, identity (4.24) implies that the objective is the solution to the equation in $h \in \mathcal{H}$,

$$T_{\mathcal{G}} \circ \mathcal{T} h = T_{\mathcal{G}} \rho. \quad (4.26)$$

On the other hand, identity (4.25) implies that the objective is the image under $T_{\mathcal{H}}^{1/2}$ of the solution to the equation in $h' \in \mathcal{L}^2(P_W)$

$$T_{\mathcal{G}}^{1/2} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2} h' = T_{\mathcal{G}}^{1/2} \rho. \quad (4.27)$$

If a solution to equation (4.26) exists in $\mathcal{H}$, then under Assumption 4.6, a solution $h_{0,KMMR}$ with minimal $\mathcal{H}$ -norm also exists and is unique. This is because, under such assumption, the map $T_{\mathcal{G}} \circ \mathcal{T} : \mathcal{H} \rightarrow \mathcal{G}$ is bounded and linear. In Appendix C.1 we show that its adjoint $(T_{\mathcal{G}} \circ \mathcal{T})^* : \mathcal{G} \rightarrow \mathcal{H}$ is given by $T_{\mathcal{H}} \circ \mathcal{T}^*$ where $\mathcal{T}^* : \mathcal{L}^2(P_Z) \rightarrow \mathcal{L}^2(P_W)$ is the adjoint of the conditional mean operator $\mathcal{T} : \mathcal{L}^2(P_W) \rightarrow \mathcal{L}^2(P_Z)$. By the identity (4.24) and Theorem 4.2 we then have that if $h_{0,KMMR}$ satisfies the source condition that posits the existence of $\beta > 0$ such that

$$h_{0,KMMR} = (T_{\mathcal{H}} \circ \mathcal{T}^* \circ T_{\mathcal{G}} \circ \mathcal{T})^{\beta/2} h^* \quad \text{for some } h^* \in \mathcal{H} \quad (4.28)$$

equivalently,

$$h_{0,KMMR} = (T_{\mathcal{H}} \circ \mathcal{T}^* \circ T_{\mathcal{G}} \circ \mathcal{T})^{\beta/2} \circ T_{\mathcal{H}}^{1/2} h^{**} \quad \text{for some } h^{**} \in \mathcal{L}^2(P_W) \quad (4.29)$$

then, there exist constants c_1 and c_2 such that

$$\|h_{0,KMMR} - h_{\lambda_{\mathcal{H},KMMR}}^{\dagger}\|_{\mathcal{H}}^2 \leq c_1 \lambda_{\mathcal{H}}^{\min\{\beta,2\}} \quad \text{and} \quad \|T_{\mathcal{G}} \circ \mathcal{T}(h_{0,KMMR} - h_{\lambda_{\mathcal{H},KMMR}}^{\dagger})\|_{\mathcal{G}}^2 \leq c_2 \lambda_{\mathcal{H}}^{\min\{\beta+1,2\}} \quad (4.30)$$

On the other hand, the identity (4.25) points out to a second possible source condition for arriving at the bounds (4.30). Specifically, consider the equation (4.27) in $h' \in \mathcal{L}^2(P_W)$. It is easy to show that the adjoint of the operator $T_G^{1/2} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2} : \mathcal{L}^2(P_W) \rightarrow \mathcal{L}^2(P_Z)$ is given by $T_{\mathcal{H}}^{1/2} \circ \mathcal{T}^* \circ T_G^{1/2} : \mathcal{L}^2(P_Z) \rightarrow \mathcal{L}^2(P_W)$. Thus, arguing as in Section 4.5, it can be shown that under the β -source condition that posits the existence of $\beta > 0$ such that

$$h_{0,KMMR} = T_{\mathcal{H}}^{1/2} (T_{\mathcal{H}}^{1/2} \circ \mathcal{T}^* \circ T_G \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2})^{\beta/2} h^{**} \quad \text{for some } h^{**} \in \mathcal{L}^2(P_W) \quad (4.31)$$

the bounds (4.30) hold for some constants c_1 and c_2 .

As in the KRAS case, Assumption (4.31) is especially convenient because it rephrases the β -source condition in terms of the operator $T_{\mathcal{H}}^{1/2} \circ \mathcal{T}^* \circ T_G \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2}$ acting between $\mathcal{L}^2$ spaces. This transformation shifts the problem from its original formulation—where the relevant operator has domain in the RKHS $\mathcal{H}$ —to one where both the loss and the penalization are expressed in $\mathcal{L}^2$ norms. In doing so, it places the KMMR estimator on the same footing as the LRAS and the KRAS estimators—the latter regarded as estimators of the image under $T_{\mathcal{H}}^{1/2}$ of the solution of the transformed equation (4.17)—enabling a direct comparison of the operators that implicitly define their associated integral equations.

Let us now compare the advantages and limitations of the KRAS and KMMR estimators.

- KRAS estimators implicitly require equation (4.17) to admit a solution in $\mathcal{L}^2(P_W)$, whereas KMMR estimators require a solution to equation (4.27). The former condition implies the latter, but not conversely—unless the integral operator T_G is injective. Injectivity holds when the kernel K_G is Matérn, but not necessarily otherwise.
- Unlike KMMR estimators, deriving bounds for KRAS estimators requires Assumption 4.5 on closeness and boundedness, which, as noted earlier, can be hard to justify.
- Even when $T_G : \mathcal{L}^2(P_Z) \rightarrow \mathcal{L}^2(P_Z)$ is injective, so that equations (4.17) and (4.27)

share the same solutions, the formulation (4.27) poses a more ill-posed problem. This is because $T_{\mathcal{G}}^{1/2} : \mathcal{L}^2(P_Z) \rightarrow \mathcal{G}$ is compact when $\mathcal{G}$ is endowed with the $\mathcal{L}^2(P_Z)$ -norm. Thus, $T_{\mathcal{G}}^{1/2}$ in equation (4.27) applies additional smoothing beyond $\mathcal{T} \circ T_{\mathcal{H}}^{1/2}$. Intuitively, solving a more ill-posed problem to a given level of accuracy demands stronger smoothness assumptions on the true solution. This added difficulty is reflected in the source condition imposed on the minimal $\mathcal{H}$ -norm solution: for the same value of β , the source condition condition (4.31) (condition (4.28)) entails stronger smoothness than the source condition in Assumption 4.4 (Assumption 4.3), because the former is formulated in terms of the fractional power of an operator that smooths more strongly than the operator appearing in the latter.

- To the best of our knowledge, the sharpest available bounds for the RMSE of $\hat{h}_{KMMR}$ under Assumption (4.28) were obtained by Park et al. (2024), refining the earlier analysis of Mastouri et al. (2021). Specifically, Park et al. (2024) showed that, with an appropriately chosen $\lambda_{\mathcal{H}}$,

$$\left\| \hat{h}_{KMMR} - h_{0,KMMR} \right\|_2 = O_p \left(n^{-\frac{\min\{\beta, 2\}}{2(1+\min\{\beta, 2\})}} \right).$$

At present, no results provide sharper bounds for the weak error $\|\mathcal{T}(\hat{h}_{KMMR} - h_{0,KMMR})\|_2$. This may be due to the fact that applying Theorem 4.2 yields bounds on the weak regularization error with respect to the transformed operator $T_{\mathcal{G}} \circ \mathcal{T}$ (as in the second inequality in (4.30) rather than $\mathcal{T}$ itself). The lack of a sharper bound on the weak error has the following consequence. Suppose we estimate the solutions to equations (4.2) and (4.3) using KMMR estimators, and that the minimal RKHS-norm solutions satisfy the β -source condition (4.28) or (4.29) with parameters β_h^{KMMR} and β_g^{KMMR} respectively. Let $\hat{\theta}_{KMMR}$ be the resulting DML estimator of θ_0 . Using the existing

finite-sample bounds we can establish that $R_n = o_p(n^{-1/2})$ if $\min\{\beta_h^{KMMR}, 2\}/(1 + \min\{\beta_h^{KMMR}, 2\} + \min\{\beta_g^{KMMR}, 2\}/(1 + \min\{\beta_g^{KMMR}, 2\})) > 1$. Thus, we cannot ensure that $R_n = o_p(n^{-1/2})$ and thus, that $\hat{\theta}_{KMMR}$ is $\sqrt{n}$ -consistent when, for instance, $\beta^{KMMR} = 1$ where $\beta^{KMMR} := \max\{\beta_h^{KMMR}, \beta_g^{KMMR}\}$.

In contrast, suppose that to estimate the solutions of equations (4.2) and (4.3) we use KRAS estimators with the RKHSs in the inner maximization for the two estimation problems having the same critical radius δ_n . Suppose that the minimum RKHS norm solutions satisfy the β -source condition of Assumption 4.3 or Assumption 4.4 with parameters β_h^{KRAS} and β_g^{KRAS} . Let $\hat{\theta}_{KRAS}$ be the DML estimator of θ_0 using the optimal choice of tuning parameters as given in Corollary 4.2. Then, using the bounds in that corollary we can establish that $R_n = o_p(n^{-1/2})$ if

$$\delta_n^{1 + \frac{\min\{\beta^{KRAS}, 1\}}{1 + \min\{\beta^{KRAS}, 1\}}} = o_p(n^{-1/2}) \quad (4.32)$$

where $\beta^{KRAS} := \max\{\beta_h^{KRAS}, \beta_g^{KRAS}\}$. As noted in Section 4.5, when the inner maximization RKHSs are Gaussian, condition (4.32) is always satisfied because $\delta_n = O(\sqrt{\{\log(n)/n\}})$. When these are d -dimensional RKHSs with Sobolev kernels of order $\nu = (l + d)/2$ where $l > d$ is a positive integer, then $\delta_n = O(n^{-\nu/(2\nu+d)})$, which yields $R_n = o_p(n^{-1/2})$ whenever $\beta^{KRAS} \in (d/l, 1]$.

A rigorous comparison of the smoothness requirements under which the remainder R_n is $o_P(n^{-1/2})$ for $\hat{\theta}_{KMMR}$ and $\hat{\theta}_{KRAS}$ is not possible for the following reasons:

1. Qualitatively, we expect that a given β^{KRAS} corresponds to a smaller β^{KMMR} for roughly the same smoothness of the minimal $\mathcal{H}$ -norm solution, though the exact correspondence between the two parameters remains unknown.
2. The absence of sharp bounds on the weak error of the KMMR nuisance parameter

estimators means that current error analysis for $\hat{\theta}_{\text{KMMR}}$ relies on bounding the product of their RMSEs. In contrast, the analysis for $\hat{\theta}_{\text{KRAS}}$ relies on bounds on the minimum of two terms, each being the product of the weak error of one nuisance estimator and the RMSE of the other. Since weak errors are typically much smaller than RMSEs, yet-unavailable sharper bounds on the weak error of the KMMR nuisance estimators could potentially lead to a substantially different error characterization for $\hat{\theta}_{\text{KMMR}}$.

3. Beyond the lack of weak error bounds, it also remains an open question whether existing RMSE bounds for KMMR estimators can be improved. This is because, the current RMSE bounds, derived by [Mastouri et al. \(2021\)](#) and [Park et al. \(2024\)](#), were obtained via an application of the Bennett–Bernstein inequality in Hilbert spaces (Theorem 6.14 of [Steinwart and Christmann \(2008\)](#)), which provides a global bound that does not account for variation in the complexity of the RKHS induced by different kernels. In contrast, following the approach of [Bennett et al. \(2023b\)](#), we instead leveraged the critical radius—a localized measure of complexity—which enabled bounds that adapt to the effective size of the RKHS. This raises the question of whether a similarly localized analysis could yield sharper convergence rates for the RMSE of KMMR estimators.
4. All that said, we conjecture that no refinement of the current analysis will show that $\hat{\theta}_{\text{KMMR}}$ estimators outperform $\hat{\theta}_{\text{KRAS}}$ estimators in terms of rates of convergence of the remainder R_n , as the fundamental issue remains: KMMR estimators target the solution of a more severely ill-posed problem, inherently limiting their precision relative to KRAS estimators.

4.7 Discussion

We have introduced a new theoretical framework for analyzing KRAS estimators arising from an auto-DML scheme. Our analysis provides finite-sample bounds on both their weak error and RMSE, under source conditions. A key strength of this framework is that it avoids reliance on hard-to-quantify measures of ill-posedness and does not require the solution to be unique. Several important directions for future research remain open.

First, our analysis has focused on single-step Tikhonov regularization. As indicated in Section 4.6.1, we conjecture that under our source conditions it will be possible to derive results for (adaptive and non-adaptive) iterated Tikhonov KRAS estimators, along the lines of [Bennett et al. \(2023b\)](#).

Second, our current analysis of KRAS estimators—and of all RAS estimators for that matter—relies on the closeness and boundedness condition in Assumption 4.5. It remains an open question whether this Assumption can be relaxed, perhaps by exploiting part 2 of Proposition 4.2.

Third, to our knowledge, the minimax optimality of LRAS estimators under source conditions has not yet been established. Likewise, it would be important to investigate if KRAS estimators are minimax optimal under our proposed source conditions.

Finally, while RAS estimators have gained popularity in recent years, the literature offers little guidance on how to tune their hyperparameters in practice. When hyperparameter selection strategies are discussed, they are often heuristic and lack theoretical guarantees. Developing principled, data-driven methods for selecting regularization parameters—and establishing their statistical properties—is a critical step toward making these estimators truly practical.

Appendix A

TECHNICAL DETAILS OF CHAPTER 2

A.1 Formulas for Unconditional and Conditional Power

Given that $\hat{\gamma}_{CP,H} \sim \mathcal{N}(\gamma_{CP,H}, V_{CP,H})$ and $\hat{\gamma}_{XC} \sim \mathcal{N}(\gamma_{XC}, V_{XC})$, the unconditional distribution of $\hat{\gamma}_{XC} + (1-f)(1+\lambda_1)\hat{\gamma}_{CP,H}$ is

$$\mathcal{N}\left(\gamma_{XC} + (1-f)(1+\lambda_1)\gamma_{CP,H}, V_{XC} + (1-f)^2(1+\lambda_1)^2V_{CP,H}\right).$$

Hence, the unconditional distribution of $T_{v,w,\lambda_1,f,\Delta_0}$ is

$$\mathcal{N}\left(\frac{\gamma_{XC} + (1-f)(1+\lambda_1)\gamma_{CP,H} - \Delta_0}{\sqrt{\tilde{V}}}, \sigma^2\right),$$

where $\sigma^2 = [V_{XC} + (1-f)^2(1+\lambda_1)^2V_{CP,H}] / \tilde{V}$ and

$$\tilde{V} = V_{XC} + 2vw(1-f)(1+\lambda_1)\sqrt{V_{XC}V_{CP,H}} + w^2(1-f)^2(1+\lambda_1)^2V_{CP,H}.$$

Moreover, at the boundary of the operational null hypothesis (2.2), the distribution of $T_{v,w,\lambda_1,f,\Delta_0}$ has mean zero for synthesis methods ($v = 0$) and mean $-Z_{1-\theta}\sqrt{\frac{V_{CP,H}}{\tilde{V}}}$ for fixed-margin methods ($v = 1$).

Next, the unconditional power of $T_{v,w,\lambda_1,f,\Delta_0}$ to reject the scientific null hypothesis (2.1) when the experimental intervention has an effect size γ_{XP} is

$$P(T_{v,w,\lambda_1,f,\Delta_0} < -Z_{1-\alpha}) = P\left\{\frac{\hat{\gamma}_{XC} + (1-f)(1+\lambda_1)\hat{\gamma}_{CP,H} - \Delta_0}{\sqrt{\tilde{V}}} < -Z_{1-\alpha}\right\}$$

$$\begin{aligned}
&= P \left\{ \hat{\gamma}_{XC} + (1-f)(1+\lambda_1)\hat{\gamma}_{CP,H} < \Delta_0 - Z_{1-\alpha}\sqrt{\tilde{V}} \right\} \\
&\stackrel{(i)}{=} P \left\{ Z < \frac{\Delta_0 - (1-f)(1+\lambda_1)\gamma_{CP,H} - \gamma_{XC} - Z_{1-\alpha}\sqrt{\tilde{V}}}{\sqrt{V_{XC} + (1-f)^2(1+\lambda_1)^2V_{CP,H}}} \right\} \\
&= \Phi \left(\frac{\Delta_0 - (1-f)(1+\lambda_1)\gamma_{CP,H} - \gamma_{XC} - Z_{1-\alpha}\sqrt{\tilde{V}}}{\sqrt{V_{XC} + (1-f)^2(1+\lambda_1)^2V_{CP,H}}} \right) \\
&\stackrel{(ii)}{=} \Phi \left(\frac{\Delta_0 + \{(1+\lambda_0) - (1-f)(1+\lambda_1)\}\gamma_{CP,H} - \gamma_{XP} - Z_{1-\alpha}\sqrt{\tilde{V}}}{\sqrt{V_{XC} + (1-f)^2(1+\lambda_1)^2V_{CP,H}}} \right),
\end{aligned}$$

where in the step (i) we first subtracted $\gamma_{XC} + (1-f)(1+\lambda_1)\gamma_{CP,H}$ and then divide by $\sqrt{V_{XC} + (1-f)^2(1+\lambda_1)^2V_{CP,H}}$ in both sides, and in (ii) we used that $\gamma_{XC} = \gamma_{XP} - \gamma_{CP} = \gamma_{XP} - (1+\lambda_0)\gamma_{CP,H}$.

In contrast, conditioning on the historical data, the conditional power of $T_{v,w,\lambda_1,f,\Delta_0}$ to reject the scientific null hypothesis (2.1) when the experimental intervention has an effect size γ_{XP} is

$$\begin{aligned}
&P(T_{v,w,\lambda_1,f,\Delta_0} < -Z_{1-\alpha}) \\
&= P \left\{ \frac{\hat{\gamma}_{XC} + (1-f)(1+\lambda_1)\hat{\gamma}_{CP,H} - \Delta_0}{\sqrt{\tilde{V}}} < -Z_{1-\alpha} \right\} \\
&= P \left\{ Z < \frac{\Delta_0 - (1-f)(1+\lambda_1)\hat{\gamma}_{CP,H} - \gamma_{XC} - Z_{1-\alpha}\sqrt{\tilde{V}}}{\sqrt{V_{XC}}} \right\} \\
&= \Phi \left(\frac{\Delta_0 - (1-f)(1+\lambda_1)\hat{\gamma}_{CP,H} - \gamma_{XC} - Z_{1-\alpha}\sqrt{\tilde{V}}}{\sqrt{V_{XC}}} \right) \\
&= \Phi \left(\frac{\Delta_0 + \{(1+\lambda_0) - (1-f)(1+\lambda_1)\}\hat{\gamma}_{CP,H} - \gamma_{XP} - Z_{1-\alpha}\sqrt{\tilde{V}}}{\sqrt{V_{XC}}} \right),
\end{aligned}$$

where in the last step, we used the relationship $\gamma_{XC} = \gamma_{XP} - \gamma_{CP} = \gamma_{XP} - (1+\lambda_0)\gamma_{CP,H}$ along with the assumption that $\gamma_{CP,H} = \hat{\gamma}_{CP,H}$. This equality is justified since we treat $\hat{\gamma}_{CP,H}$ as fixed and its mean is equal to $\gamma_{CP,H}$.

For fixed-margin methods, where $v = 0$, $w = \frac{Z_{1-\theta}}{Z_{1-\alpha}}$, and $\lambda_1 = 0$, it follows that $\sqrt{\tilde{V}} = \sqrt{V_{XC}} + \frac{Z_{1-\theta}}{Z_{1-\alpha}}(1-f)\sqrt{V_{CP,H}}$. Hence, the previous expression can be simplified as follows:

$$\begin{aligned} & \Phi \left(\frac{\Delta_0 - (1-f)\hat{\gamma}_{CP,H} - \gamma_{XC} - Z_{1-\alpha}\sqrt{\tilde{V}}}{\sqrt{V_{XC}}} \right) \\ &= \Phi \left(\frac{\Delta_0 - (1-f)\{\hat{\gamma}_{CP,H} + Z_{1-\theta}\sqrt{V_{CP,H}}\} - \gamma_{XC}}{\sqrt{V_{XC}}} - Z_{1-\alpha} \right) \\ &= \Phi \left(-Z_{1-\alpha} + \frac{\delta - \gamma_{XC}}{\sqrt{V_{XC}}} \right), \end{aligned}$$

where δ is the success margin defined in (2.11). This expression is commonly used to compute the required precision, V_{XC} , of the relative effect of experimental versus active control intervention that ensures adequate conditional power and type I error control under constancy.

A.2 Detectable Design Alternatives and Maximum Unconditional Power

Define $x := \sqrt{V_{XC}}$, $a = \Delta_0 + \{(1 + \lambda_0) - (1 - f)(1 + \lambda_1)\}\gamma_{CP,H} - \gamma_{XP}$, and $\varepsilon := (1 - f)(1 + \lambda_1)\sqrt{V_{CP,H}}$. For $\varepsilon > 0$, $x > 0$, and $v, w \geq 0$, it follows that $w\varepsilon < \sqrt{x^2 + 2vw\varepsilon x + w^2\varepsilon^2}$. Thus, if $a - w\varepsilon Z_{1-\alpha} \leq 0$, we have

$$a - Z_{1-\alpha}\sqrt{x^2 + 2vw\varepsilon x + w^2\varepsilon^2} < a - w\varepsilon Z_{1-\alpha} \leq 0 \quad \forall x > 0,$$

implying that both conditional and unconditional powers are less than 0.5. Conversely, for any $v \in [0, 1]$ and $x > 0$, the following holds:

$$\begin{aligned} \frac{a - Z_{1-\alpha}\sqrt{x^2 + 2vw\varepsilon x + w^2\varepsilon^2}}{x} &= \frac{a}{x} - Z_{1-\alpha}\sqrt{1 + \frac{2vw\varepsilon}{x} + \frac{w^2\varepsilon^2}{x^2}} \\ &\geq \frac{a}{x} - Z_{1-\alpha}\left(1 + \frac{w\varepsilon}{x}\right) \\ &= -Z_{1-\alpha} + \frac{a - w\varepsilon Z_{1-\alpha}}{x}. \end{aligned}$$

Therefore, if $a - w\varepsilon Z_{1-\alpha} > 0$, this expression can be made arbitrarily positive by an appropriate choice of x . Consequently, conditional power can exceed 0.5, meaning that to detect a design alternative with 50% or greater conditional power, the inequality $a - w\varepsilon Z_{1-\alpha} > 0$ must hold. This condition is equivalent to γ_{XP} being less than

$$\Delta_0 + (1 + \lambda_0)\gamma_{CP,H} - (1 - f)(1 + \lambda_1) \left\{ \gamma_{CP,H} + wZ_{1-\alpha}\sqrt{V_{CP,H}} \right\}.$$

Additionally, for any $x > 0$, it holds that $\varepsilon < \sqrt{x^2 + \varepsilon^2}$, which implies that

$$\frac{a - Z_{1-\alpha}\sqrt{x^2 + 2vw\varepsilon x + w^2\varepsilon^2}}{\sqrt{x^2 + \varepsilon^2}} < \frac{a - w\varepsilon Z_{1-\alpha}}{\varepsilon}.$$

Thus, the unconditional power is bounded above by $\Phi\left(\frac{a - w\varepsilon Z_{1-\alpha}}{\varepsilon}\right)$. Expressing this bound in terms of the original parameters yields

$$\Phi\left(-wZ_{1-\alpha} + \frac{\Delta_0 + \{(1 + \lambda_0) - (1 - f)(1 + \lambda_1)\}\gamma_{CP,H} - \gamma_{XP}}{(1 - f)(1 + \lambda_1)\sqrt{V_{CP,H}}}\right).$$

Finally, to detect a design alternative with an unconditional power of $(1 - \beta)$, the following condition must be satisfied:

$$\frac{a - w\varepsilon Z_{1-\alpha}}{\varepsilon} > Z_{1-\beta},$$

which is equivalent to γ_{XP} being less than

$$\Delta_0 + (1 + \lambda_0)\gamma_{CP,H} - (1 - f)(1 + \lambda_1) \left\{ \gamma_{CP,H} + (wZ_{1-\alpha} + Z_{1-\beta})\sqrt{V_{CP,H}} \right\}.$$

A.3 Acceptable Assumed Level of Non-Constancy and Tolerable Level of Non-Constancy

To ensure that the unconditional type I error remains below 0.5, the numerator in expression (2.6) must be negative. This leads to the requirement $(1-f)(\lambda_0 - \lambda_1)\gamma_{CP,H} - Z_{1-\alpha}\sqrt{\tilde{V}} < 0$. Given that V_{XC} could take any positive value, we want the inequality above to hold for all $V_{XC} > 0$. This condition is satisfied if

$$(1-f)(\lambda_0 - \lambda_1)\gamma_{CP,H} \leq w(1 + \lambda_1)(1-f)Z_{1-\alpha}\sqrt{V_{CP,H}} \quad (\text{A.1})$$

If $\lambda_1 \leq \lambda_0$, the inequality (A.1) always hold. On the other hand, when $\lambda_1 > \lambda_0$, assuming that $\gamma_{CP,H} + wZ_{1-\alpha}\sqrt{V_{CP,H}} < 0$, (A.1) holds if

$$\lambda_1 \leq \lambda_0 - \frac{(1 + \lambda_0)wZ_{1-\alpha}\sqrt{V_{CP,H}}}{\gamma_{CP,H} + wZ_{1-\alpha}\sqrt{V_{CP,H}}}.$$

The assumption that $\gamma_{CP,H} + wZ_{1-\alpha}\sqrt{V_{CP,H}} < 0$ is reasonable given that the active control was effective in the historical setting. Specifically, if $\gamma_{CP,H} = \hat{\gamma}_{CP,H}$, then $\gamma_{CP,H} + wZ_{1-\alpha}\sqrt{V_{CP,H}} < 0$ represents the upper bound of a confidence interval for $\gamma_{CP,H}$, which we expect to lie below zero.

To quantify the level of non-constancy that a given method in the general framework properly accommodates, we solve the following inequality for λ_0 :

$$\frac{(1-f)(\lambda_0 - \lambda_1)\gamma_{CP,H} - Z_{1-\alpha}\sqrt{\tilde{V}}}{\sqrt{V_{XC} + (1-f)^2(1 + \lambda_1)^2 V_{CP,H}}} \leq -Z_{1-\alpha}.$$

Given that $\gamma_{CP,H} < 0$, the solution to this inequality yields $\lambda_0 \geq \lambda_{0,\min}$, where

$$\lambda_{0,\min} := \lambda_1 + \frac{\sqrt{\tilde{V}} - \sqrt{V_{XC} + (1-f)^2(1 + \lambda_1)^2 V_{CP,H}}}{(1-f)\gamma_{CP,H}} Z_{1-\alpha}.$$

This value, $\lambda_{0,\min}$, defines the tolerable level of non-constancy.

For the bias-adjusted synthesis method, $\sqrt{\tilde{V}} = \sqrt{V_{XC} + (1-f)^2(1+\lambda_1)^2V_{CP,H}}$. In this case, $\lambda_{0,\min} = \lambda_1$. For the Odem-Davis method, we find that

$$\sqrt{\tilde{V}} = \sqrt{V_{XC} + (1-f)^2V_{CP,H}} > \sqrt{V_{XC} + (1-f)^2(1+\lambda_1)^2V_{CP,H}}$$

for any $\lambda_1 \in (-1, 0)$. Therefore, since $\gamma_{CP,H}$ is negative, it follows that $\lambda_{0,\min} < \lambda_1$.

For the fixed-margin method, where $\lambda_1 = 0$ and $\sqrt{\tilde{V}} = \sqrt{V_{XC}} + \frac{Z_{1-\theta}}{Z_{1-\alpha}}(1-f)\sqrt{V_{CP,H}}$, we obtain

$$\lambda_{0,\min} = \frac{\sqrt{V_{XC}} + \frac{Z_{1-\theta}}{Z_{1-\alpha}}(1-f)\sqrt{V_{CP,H}} - \sqrt{V_{XC} + (1-f)^2V_{CP,H}}}{(1-f)\gamma_{CP,H}}Z_{1-\alpha}. \quad (\text{A.2})$$

Since $\gamma_{CP,H}$ is assumed to be negative, the right hand side of (A.2) is negative for the 95-95 method (with $\theta = \alpha = 0.025$), indicating that $\lambda_{0,\min} < 0$. In contrast, for the 0-95 method (where $Z_{1-\theta} = 0$), the right hand side of (A.2) becomes positive, resulting in $\lambda_{0,\min} > 0$.

Appendix B

SUPPLEMENTARY MATERIAL OF CHAPTER 3

B.1 Proximal Identification and Estimation in a potentially restricted population

Throughout we let P denote the distribution of $O = (L, Z, W, A, Y)$. The symbol $\mathbb{E}$ stands for expectation under P . In this Appendix we extend the identifiability results and estimators in Chapter 3 to target estimands of the form

$$\psi_0 = \mathbb{E} [Y \{q(A, L)\} | (A, L) \in \mathcal{S}], \quad (\text{B.1})$$

where $q : \text{supp}(A, L) \mapsto \mathbb{R}$ is a specified map and $\mathcal{S}$ is a specified subset of $\text{supp}(A, L)$ such that $P\{(A, L) \in \mathcal{S}\} > 0$ and that for every $(a, l) \in \mathcal{S}$, $(q(a, l), l)$ belongs to $\text{supp}(A, L)$. Of course, the estimands studied in Chapter 3 are also of the form (B.1) but with $\mathcal{S} = \text{supp}(A, L)$.

Note that the target estimand in Example 1 of Chapter 3 is of this form where $q(a, l) = a + \delta$ and $\mathcal{S} = \{(a, l) : l \in \text{supp}(L), a \in [c(l), d(l) - \delta]\}$. However, subpopulations $\{O : (A, L) \in \mathcal{S}'\}$, where $\mathcal{S}'$ is a strict subset of $\mathcal{S}$, may also be of interest, particularly for comparing the counterfactual MTP mean across different policies within a common population $\{O : (A, L) \in \mathcal{S}'\}$.

B.1.1 Proximal Identifying Assumptions

For each $l \in \text{supp}(L)$, define $\mathcal{S}(l) := \{a \in \text{supp}(A) : (a, l) \in \mathcal{S}\}$.

Assumption B.1 (Consistency). $Y(A) = Y$.

Assumption B.2 (Randomization of policy related populations). *For all $(a, l, u) \in \text{supp}(A, L, U)$ such that $x \in \mathcal{S}(l)$ and $a' = q(a, l)$, the conditional distribution of $Y(a')|A = a, L = l, U = u$ is equal to the conditional distribution of $Y(a')|A = a', L = l, U = u$.*

Assumption B.3 (Common support). *For all $(l, u) \in \text{supp}(L, U)$, $\text{supp}(A|L = l, U = u) = \text{supp}(A|L = l)$.*

Assumption B.4 (Negative control treatment). $Z \perp Y|A, L, U$.

Assumption B.5 (Negative control outcome). $(A, Z) \perp W|U, L$.

Assumption B.6 (Monotone and smooth policy). *For each $l \in \text{supp}(L)$, the map $q(\cdot, l) : \mathcal{S}(l) \mapsto \text{supp}(A)$ is strictly monotone and differentiable almost everywhere with respect to the distribution of A given $L = l$.*

Assumption B.6 ensures that, for each $l \in \text{supp}(L)$, the policy map $q : a \mapsto q(a, l)$ has an inverse $q^{-1} : q(\mathcal{S}(l), l) \mapsto \mathcal{S}(l)$ that is well-defined and almost everywhere differentiable with respect to the distribution of A given $L = l$.

Assumption B.7 (Existence of latent outcome bridge and observed treatment bridge functions).

- i) There exists $h_0 \in \mathcal{L}^2(A, L, W)$, hereafter called “latent outcome bridge function”, solving for all $(a, l, u) \in \text{supp}(A, L, U)$ such that $a \in \mathcal{S}(l) \cup q(\mathcal{S}(l), l)$, the integral equation*

$$\mathbb{E}(Y|A = a, L = l, U = u) = \int h(a, l, w) p_{W|A, L, U}(w|a, l, u) dw. \quad (\text{B.2})$$

- ii) There exists $g_0 \in \mathcal{L}^2(A, L, Z)$, hereafter called “observed treatment bridge function”, satisfying $g_0(a, l, z) = 0$ if $a \in \mathcal{S}(l) - q(\mathcal{S}(l), l)$ and such that at every $a \in q(\mathcal{S}(l), l)$ at*

which the map $a' \mapsto q^{-1}(a', l)$ is differentiable at $a' = a$, g_0 solves the integral equation

$$\alpha_0(a, l, w) = \int g(a, l, z) p_{Z|A, L, W}(z|a, l, w) dz, \quad (\text{B.3})$$

where

$$\alpha_0(a, l, w) := I_q(a, l) \cdot \frac{dq^{-1}(a, l)}{da} \frac{p_{A|L, W}\{q^{-1}(a, l)|l, w\}}{p_{A|L, W}(a|l, w)}$$

and $I_q(a, l) := \{a \in q(\mathcal{S}(l), l)\}$.

Assumption B.8 (Existence of latent treatment bridge and observed outcome bridge functions).

i) There exists $g_0 \in \mathcal{L}^2(A, L, Z)$, hereafter called “latent treatment bridge function”, satisfying $g_0(a, l, z) = 0$ if $a \in \mathcal{S}(l) - q(\mathcal{S}(l), l)$ and such that at every $a \in q(\mathcal{S}(l), l)$ at which the map $a' \mapsto q^{-1}(a', l)$ is differentiable at $a' = a$, g_0 solves the integral equation

$$\alpha_0(a, l, u) = \int g(a, l, z) p_{Z|A, L, U}(z|a, l, u) dz, \quad (\text{B.4})$$

where

$$\alpha_0(a, l, u) := I_q(a, l) \frac{dq^{-1}(a, l)}{da} \frac{p_{A|L, U}\{q^{-1}(a, l)|l, u\}}{p_{A|L, U}(a|l, u)}$$

ii) There exists $h_0 \in \mathcal{L}^2(A, L, W)$, hereafter called “observed outcome bridge function”, solving for all $(a, l, w) \in \text{supp}(A, L, W)$ such that $a \in \mathcal{S}(l) \cup q(\mathcal{S}(l), l)$, the integral equation

$$\mathbb{E}(Y|A = a, L = l, Z = z) = \int h(a, l, w) p_{W|A, L, Z}(w|a, l, z) dw. \quad (\text{B.5})$$

B.1.2 Conditions for Existence of the Latent Bridge Functions

We consider first sufficient conditions for the existence of solutions to equations (B.2) and (B.4) when the variables Z , W , and U are discrete and take values in finite sets. For equation (B.2), sufficient conditions were discussed in Miao et al. (2018); Tchetgen et al. (2020), and Shi et al. (2020). For equation (B.4), conditions can be derived by mimicking the analysis given in Cui et al. (2024) and Kallus et al. (2021) for the existence of solutions to the treatment bridge equations corresponding to the average treatment effect (ATE) and generalized ATE respectively. Here, we review these analyses.

Suppose that W , Z , and U take a finite number of values, denoted by w_i , z_j , and u_k for $i = 1, \dots, n_W$, $j = 1, \dots, n_Z$, and $k = 1, \dots, n_U$, respectively. For any $(a, l, u) \in \text{supp}(A, L, U)$ such that $a \in \mathcal{S}(l) \cup q(\mathcal{S}(l), l)$, let $\mathbb{P}(\mathbf{W}|\mathbf{U}, a, l)$ denote the $n_W \times n_U$ matrix whose (i, k) -entry is $P(W = w_i | U = u_k, A = a, L = l)$. Similarly, let $\mathbb{P}(\mathbf{Z}|\mathbf{U}, a, l)$ be the $n_Z \times n_U$ matrix with (j, k) -entry $P(Z = z_j | U = u_k, A = a, L = l)$. Define $r_0(a, l, \mathbf{U})$ and $a_0(a, l, \mathbf{U})$ as the n_U -dimensional column vectors whose k -th entries are $\mathbb{E}(Y | A = a, L = l, U = u_k)$ and $\alpha_0(a, l, u_k)$, respectively. With this notation, the outcome bridge equation (B.2) is:

$$r_0(a, l, \mathbf{U}) = \mathbb{P}(\mathbf{W}|\mathbf{U}, a, l)^T h(a, l, \mathbf{W}),$$

and the treatment bridge equation (B.4) is

$$a_0(a, l, \mathbf{U}) = \mathbb{P}(\mathbf{Z}|\mathbf{U}, a, l)^T g(a, l, \mathbf{Z}).$$

It follows that a solution to equation (B.2) exists if only if $r_0(a, l, \mathbf{U})$ lies in the vector space spanned by the row vectors of $\mathbb{P}(\mathbf{W}|\mathbf{U}, a, l)$. Likewise, a solution to equation (B.4) exists if and $a_0(a, l, \mathbf{U})$ is in the vector space spanned by the rows of $\mathbb{P}(\mathbf{Z}|\mathbf{U}, a, l)$. If $n_W \geq$

n_U ($n_Z \geq n_U$) and the matrix $\mathbb{P}(\mathbf{W}|\mathbf{U}, a, l)$ ($\mathbb{P}(\mathbf{Z}|\mathbf{U}, a, l)$) is of full row order, then the solution to the outcome bridge equation (B.2) (treatment bridge equation (B.4)) always exists. If $n_W > n_U$ then equation (B.2) may have many solutions. Likewise, if $n_Z > n_U$ then equation (B.4) may have many solutions.

When W , Z , and U are continuous, as noted in Miao et al. (2018) and Cui et al. (2024), we can deduce sufficient conditions for the existence of solutions to equations (B.2) by invoking Picard's theorem (Theorem 15.18 in Kress (2014)). Specifically, equation (B.2) can be written as

$$\mathcal{T}_{a,l}^W h(a, l, U) = \mathbb{E}(Y|A = a, L = l, U), \quad (\text{B.6})$$

where for each $(a, l) \in \text{supp}(A, L)$, we let $\mathcal{T}_{a,l}^W : \mathcal{L}^2(W|A = a, L = l) \mapsto \mathcal{L}^2(U|A = a, L = l)$ denote the conditional expectation operator defined by $\mathcal{T}_{a,l}^W h(a, l, U) := \mathbb{E}\{h(A, L, W)|A = a, L = l, U\}$ and throughout we assume that $\mathbb{E}(Y|A = a, L = l, U) \in \mathcal{L}^2(U|A = a, L = l)$. Then, equation (B.6) has a solution if and only if $\mathbb{E}(Y|A = a, L = l, U)$ lies in the range of the, continuous and linear, operator $\mathcal{T}_{a,l}^W$. Now, the closure of the range of a continuous and linear operator is equal to the orthogonal complement of the null space of the adjoint of the operator. For $\mathcal{T}_{a,l}^W$, its adjoint is the operator $\mathcal{T}_{a,l}^U : \mathcal{L}^2(U|A = a, L = l) \mapsto \mathcal{L}^2(W|A = a, L = l)$ defined as $\mathcal{T}_{a,l}^U g(a, l, w) := \mathbb{E}\{g(A, L, U)|A = a, L = l, W = w\}$, whose kernel is $\{\eta(a, l, U) : \mathbb{E}[\eta(A, L, U)|A = a, L = l, W] = 0\}$. Thus, in particular, the closure of the range of the operator $\mathcal{T}_{a,l}^W$ equals $\mathcal{L}^2(U|A = a, L = l)$ if and only if kernel of $\mathcal{T}_{a,l}^U$ is the set $\{0\}$, equivalently, if and only if

$$\eta(a, l, U = u) = 0 \quad \forall u \quad \text{whenever} \quad \mathbb{E}\{\eta(A, L, U)|A = a, L = l, W = w\} = 0 \quad \forall w. \quad (\text{B.7})$$

Condition (B.7) is the completeness condition discussed in Chapter 3. So, we conclude that the completeness condition ensures that the closure of the range of the operator $\mathcal{T}_{a,l}^W$

is the entire space $\mathcal{L}^2(U|A = a, L = l)$. However, this is not sufficient to ensure that $\mathbb{E}\{Y|A = a, L = l, U\}$ is in the range of $\mathcal{T}_{a,l}^W$, as it may fall in the boundary of the range. Picard's Theorem provides a sufficient condition for a given element of $\mathcal{L}^2(U|A = a, L = l)$, in our case $\mathbb{E}\{Y|A = a, L = l, U\}$, to not fall in the boundary, provided the operator $\mathcal{T}_{a,l}^W$ is compact. Compactness of $\mathcal{T}_{a,l}^W$ holds under the regularity condition that (see Examples 2.2 and 2.3 in Carrasco et al. (2007)):

$$\iint p_{W|A,L,U}(w|a, l, u)p_{U|A,L,W}(u|a, l, w)dwdu < \infty. \quad (\text{B.8})$$

Under (B.8), the operator $\mathcal{T}_{a,l}^W$ admits a singular value decomposition $(\sigma_{a,l,j}, \varphi_{a,l,j}, \phi_{a,l,j})_{j=1}^{\infty}$ (see Theorem 15.16 in Kress (2014)). Under completeness and compactness, Piccard's sufficient condition for $\mathbb{E}\{Y|A = a, L = l, U\}$ to fall outside the boundary of the range of $\mathcal{T}_{a,l}^W$ is

$$\sum_{j=1}^{\infty} \sigma_{a,l,j}^{-2} \left\{ \int \mathbb{E}(Y|A = a, L = l, U = u) \phi_{a,l,j}(a, l, u) p_{U|A,L}(u|a, l) du \right\}^2 < \infty. \quad (\text{B.9})$$

Thus, we conclude that the completeness condition (B.7) together with the regularity conditions (B.8) and (B.9) suffice for equation (B.6) to have a solution.

Likewise, we can invoke Picard's theorem to derive sufficient conditions for the existence of a solution to equation (3.3). These are the completeness condition

$$\eta(a, l, U = u) = 0 \quad \forall u \quad \text{whenever} \quad \mathbb{E}\{\eta(A, L, U)|A = a, L = l, Z = z\} = 0 \quad \forall z$$

and the regularity conditions

$$\iint p_{Z|A,L,U}(Z|a, l, u)p_{U|A,L,Z}(u|a, l, z)dzdu < \infty,$$

and

$$\sum_{j=1}^{\infty} \tilde{\sigma}_{a,l,j}^{-2} \left\{ \int \alpha_0(a, l, u) \tilde{\phi}_{a,l,j}(a, l, u) p_{U|A,L}(u|a, l) du \right\}^2 < \infty,$$

where $(\tilde{\sigma}_{a,l,j}, \tilde{\varphi}_{a,l,j}, \tilde{\phi}_{a,l,j})_{j=1}^{\infty}$ is the singular value decomposition of the compact operator $\mathcal{T}_{a,l}^Z : \mathcal{L}^2(Z|A = a, L = l) \mapsto \mathcal{L}^2(U|A = a, L = l)$ defined as

$$\mathcal{T}_{a,l}^Z g(a, l, U) := \mathbb{E} \{g(A, L, Z) | A = a, L = l, U\}.$$

B.1.3 Proximal Identification Results

Next, we extend the identification result in Chapter 3 for the counterfactual MTP mean within the population $\{O : (A, L) \in \mathcal{S}\}$. Define

$$\phi(O; h, g) := h \{q(A, L), L, W\} \cdot I \{(A, L) \in \mathcal{S}\} + g(A, L, Z) \{Y - h(A, L, W)\}. \quad (\text{B.10})$$

Theorem B.1. *Suppose Assumptions B.1–B.6 hold. If, in addition, either Assumption B.7 holds or Assumption B.8 holds, then ψ_0 is identified by the observed data and it satisfies all of the following:*

$$\psi_0 = \frac{\mathbb{E} [h_0 \{q(A, L), L, W\} \cdot I \{(A, L) \in \mathcal{S}\}]}{P \{(A, L) \in \mathcal{S}\}} \quad (\text{B.11})$$

$$\begin{aligned} &= \frac{\mathbb{E} \{Y g_0(A, L, Z)\}}{P \{(A, L) \in \mathcal{S}\}} \\ &= \frac{\mathbb{E} \{\phi(O; h^\dagger, g^\dagger)\}}{P \{(A, L) \in \mathcal{S}\}} \end{aligned} \quad (\text{B.12})$$

where h_0 is an observed outcome bridge function, g_0 is an observed treatment bridge function and, either $h^\dagger$ is an observed outcome bridge function or $g^\dagger$ is an observed treatment bridge function but not necessarily both functions are observed bridge functions.

Theorem B.1 offers a slight weakening of the conditions for the outcome regression representation (B.11) of the preceding theorem that would follow if one were to apply the identification strategy in Kallus et al. (2021) to our target parameter ψ_0 . Specifically, following the identification analysis in Kallus et al. (2021) we arrive at the conclusion that the representation holds under Assumption B.7. Theorem B.1 establishes that this representation is valid even when Assumption B.7 fails as long as Assumption B.8 holds. Likewise, for the representation (B.12), Theorem B.1 establishes its validity even if Assumption B.8 fails, provided that Assumption B.7 holds.

B.1.4 Estimation

The estimator $\hat{\psi}_{CF}^{DR}$ of Chapter 3 extends to the following estimator for the target parameter defined in (B.1) of this Appendix:

$$\hat{\psi}_{CF}^{DR} := \frac{\sum_{k=1}^K \sum_{i: O_i \in I_k} \phi \{O_i; \hat{h}^{(-k)}, \hat{g}^{(-k)}\}}{\sum_{i=1}^n I \{(A_i, L_i) \in \mathcal{S}\}}.$$

where $\hat{h}^{(-k)}$ and $\hat{g}^{(-k)}$ are defined as in Chapter 3 except that their computation is restricted to the units in all the sample but sample I_k and that satisfy that $(A, L) \in \mathcal{S}$.

Note that when $\mathcal{S} = \text{supp}(A, L)$, i.e., when $I \{(A, L) \in \mathcal{S}\} \equiv 1$, this estimator $\hat{\psi}_{CF}^{DR}$ coincides with the estimator proposed in Chapter 3.

The following theorem extends the Theorem 3.3 in Chapter 3.

Theorem B.2. *Suppose that the Assumptions of Theorem B.1 hold. Let h_0 and g_0 denote observed outcome and treatment bridge functions respectively. If for some $B > 0$, $|\alpha_0(A, L, W)| \leq B$, $|Y| \leq B$ and for each k , $\|\hat{h}^{(-k)} - h_0\|_2 = o_p(1)$, $\|\hat{g}^{(-k)} - g_0\|_2 = o_p(1)$ and either $\|h_0\|_\infty + \|\hat{g}^{(-k)}\|_\infty \leq B$ or $\|\hat{h}^{(-k)}\|_\infty + \|g_0\|_\infty \leq B$ for some B , then the estimator*

$\hat{\psi}_{CF}^{DR}$ satisfies

$$\sqrt{n} \left\{ \hat{\psi}_{CF}^{DR} - \psi_0 \right\} = \sqrt{n} \mathbb{E}_n \left[p_0^{-1} \phi(O; h_0, g_0) - \psi_0 p_0^{-1} I \{ (A, L) \in \mathcal{S} \} \right] + \sqrt{n} R_n + o_p(1).$$

where $R_n = p_0^{-1} \sum_{k=1}^K \frac{n_k}{n} \int \{ \hat{h}^{(-k)}(a, l, w) - h_0(a, l, w) \} \{ g_0(a, l, z) - \hat{g}^{(-k)}(a, l, z) \} dP(a, l, w, z)$ with $p_0 := P \{ (A, L) \in \mathcal{S} \}$ and n_k denoting the cardinality of I_k . In particular, if $R_n = o_p(n^{-1/2})$ and $\mathbb{E} \{ \phi(O; h_0, g_0)^2 \} < \infty$, then

$$\sqrt{n} \left\{ \hat{\psi}_{CF}^{DR} - \psi_0 \right\} \xrightarrow{d} \mathcal{N}(0, \tau^2),$$

where $\tau^2 = \mathbb{E} \left\{ \left[\frac{\phi(O; h_0, g_0)}{P \{ (A, L) \in \mathcal{S} \}} - \psi_0 \frac{I \{ (A, L) \in \mathcal{S} \}}{P \{ (A, L) \in \mathcal{S} \}} \right]^2 \right\}$.

B.2 Estimation of the Bridge Functions

The estimators for the bridge functions introduced in Section 3.4.1 correspond to the kernel regularized adversarial stabilized (KRAS) estimators proposed by Dikkala et al. (2020); Ghassami et al. (2022), and further studied in Chapter 4. We assume that the observed bridge equations (B.5) and (B.3) admit solutions within the RKHSs used for estimation, and let h_0 and g_0 denote the corresponding solutions with minimum RKHS norm. Under this assumption, and under additional regularity conditions on the RKHSs—most of which are automatically satisfied by the RKHSs we use—Theorems 4.4 and 4.5 of Chapter 4 establish upper bounds on the root-mean-squared errors $\|\hat{h} - h_0\|_2$ and $\|\hat{g} - g_0\|_2$, as well as on the weak errors $\|\mathbb{E}\{\hat{h}(A, L, W) - h_0(A, L, W) | A, L, Z\}\|_2$ and $\|\mathbb{E}\{\hat{g}(A, L, Z) - g_0(A, L, Z) | A, L, W\}\|_2$. These rates hold under a source condition that imposes a smoothness constraint on h_0 and g_0 tailored to the geometry of the RKHSs $\mathcal{H}$ and $\mathcal{G}$. The only assumption not automatically guaranteed by our setting—beyond the smoothness conditions imposed on h_0 and g_0 —is the local closedness and boundedness condition on the RKHSs involved in the inner and

outer maximization problems, which ensures the conditional expectation operator behaves regularly when restricted to bounded subsets. We make this assumption here to apply the theoretical results from Chapter 4. See Chapter 4 for further discussion.

In particular, when h_0 and g_0 belong to Gaussian RKHS or to Sobolev RKHSs of sufficiently high order, the corresponding estimators $\hat{h}$ and $\hat{g}$ are norm-consistent and attain convergence rates such that the product of the weak error of one estimator and the root-mean-squared error of the other is $o_p(n^{-1/2})$, thereby satisfying the conditions of Theorem B.2.

B.2.1 Estimators Using Reproducing Kernel Hilbert Spaces

In this section, we derive closed-form expressions for the estimators of the bridge functions using reproducing kernel Hilbert spaces. Recall that the estimator $\hat{h}$ for h_0 is defined as the solution to the regularized minimax problem

$$\hat{h} = \arg \min_{h \in \mathcal{H}_B} \max_{g \in \mathcal{G}'} \mathbb{E}_n [g(A, L, Z) \{Y - h(A, L, W)\} - g(A, L, Z)^2] - \lambda_{\mathcal{G}'} \|g\|_{\mathcal{G}'} + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2.$$

Importantly, to satisfy the assumptions required by Theorems 4.4 and 4.5 of Chapter 4, we restrict the outer optimization to the class $\mathcal{H}_B := \{h \in \mathcal{H} : \|h\|_{\mathcal{H}} \leq B\}$, where B is a sufficiently large constant chosen so that $\|h_0\|_{\mathcal{H}} < B$.

We begin by analyzing the inner maximization. Specifically, for any function $h \in \mathcal{H}$, we show that the inner problem satisfies

$$\begin{aligned} & \max_{g \in \mathcal{G}'} \mathbb{E}_n [g(A, L, Z) \{Y - h(A, L, W)\} - g(A, L, Z)^2] - \lambda_{\mathcal{G}'} \|g\|_{\mathcal{G}'}^2 \\ &= \frac{1}{4} \{\xi_n(h)\}^T \left[\frac{1}{n} K_{\mathcal{G}', n} + \lambda_{\mathcal{G}'} I_n \right]^{-1} K_{\mathcal{G}', n} \xi_n(h), \end{aligned} \quad (\text{B.13})$$

where $K_{\mathcal{G}', n}$ is the $n \times n$ matrix with (i, j) -th entry given by $K_{\mathcal{G}'} \{(A_i, L_i, Z_i), (A_j, L_j, Z_j)\}$,

$$\xi_n(h) := \frac{1}{n} [\widetilde{Y}_n - \tilde{h}_n], \quad \widetilde{Y}_n = (Y_1, \dots, Y_n)^T \text{ and } \tilde{h}_n = [h(A_1, L_1, W_1), \dots, h(A_n, L_n, W_n)]^T.$$

By the generalized representer theorem (see [Schölkopf et al. \(2001\)](#)), the solution to the maximization problem (B.13) has the form

$$g(a, l, z) = \sum_{j=1}^n \alpha_j K_{\mathcal{G}'}\{(a, l, z), (A_j, L_j, Z_j)\},$$

where $\alpha = (\alpha_1, \dots, \alpha_n)^T \in \mathbb{R}^n$. This result ensures that the solution is fully characterized by the observed data through the kernel evaluations. For $i = 1, \dots, n$, we can write $g(A_i, L_i, Z_i)$ as:

$$g(A_i, L_i, Z_i) = e_i^T K_{\mathcal{G},n} \alpha$$

where e_i is the i -th standard basis vector in $\mathbb{R}^n$. Using this representation, the terms in (B.13) evaluated at functions g as in the last display, can be written as follows:

$$\begin{aligned} \mathbb{E}_n [Y g(A, L, Z)] &= \frac{1}{n} \sum_{i=1}^n Y_i e_i^T K_{\mathcal{G}',n} \alpha = \frac{1}{n} \widetilde{Y}_n^T K_{\mathcal{G}',n} \alpha, \\ \mathbb{E}_n [g(A, L, Z) h(A, L, W)] &= \frac{1}{n} \sum_{i=1}^n h(X_i, L_i, W_i) e_i^T K_{\mathcal{G}',n} \alpha = \frac{1}{n} \tilde{h}_n^T K_{\mathcal{G}',n} \alpha, \\ \mathbb{E}_n [g(A, L, Z)^2] &= \frac{1}{n} \sum_{i=1}^n [\alpha^T K_{\mathcal{G}',n} e_i e_i^T K_{\mathcal{G}',n} \alpha] = \frac{1}{n} \alpha^T K_{\mathcal{G}',n}^2 \alpha, \\ \|g\|_{\mathcal{G}'}^2 &= \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j K_{\mathcal{G}'}\{(A_i, L_i, Z_i), (A_j, L_j, Z_j)\} = \alpha^T K_{\mathcal{G}',n} \alpha. \end{aligned}$$

Substituting these expressions into (B.13), we obtain

$$\begin{aligned} &\mathbb{E}_n [g(A, L, Z) \{Y - h(A, L, W)\} - g(A, L, Z)^2] - \lambda_{\mathcal{G}'} \|g\|_{\mathcal{G}'}^2 \\ &= \frac{1}{n} \alpha^T K_{\mathcal{G}',n} [\widetilde{Y}_n - \tilde{h}_n] - \alpha^T \left[\frac{1}{n} K_{\mathcal{G}',n}^2 + \lambda_{\mathcal{G}'} K_{\mathcal{G}',n} \right] \alpha. \end{aligned}$$

This is a concave quadratic function of α with a unique maximizer, equal to

$$\alpha^* = \frac{1}{2n} \left[\frac{1}{n} K_{\mathcal{G}',n} + \lambda_{\mathcal{G}'} I_n \right]^{-1} [\widetilde{Y}_n - \tilde{h}_n].$$

which can be derived by taking the derivative of the quadratic function with respect to α and setting it to zero. We then arrive at

$$\begin{aligned} & \max_{g \in \mathcal{G}'} \mathbb{E}_n [g(A, L, Z) \{Y - h(A, L, W)\} - g(A, L, Z)^2] - \lambda_{\mathcal{G}'} \|g\|_{\mathcal{G}'}^2 \\ &= \frac{1}{2n^2} [\widetilde{Y}_n - \tilde{h}_n]^T \left[\frac{1}{n} K_{\mathcal{G}',n} + \lambda_{\mathcal{G}'} I_n \right]^{-1} K_{\mathcal{G}',n} [\widetilde{Y}_n - \tilde{h}_n] \\ &\quad - \frac{1}{4n^2} \left\{ \left[\frac{1}{n} K_{\mathcal{G}',n} + \lambda_{\mathcal{G}'} I_n \right]^{-1} [\widetilde{Y}_n - \tilde{h}_n] \right\}^T K_{\mathcal{G}',n} [\widetilde{Y}_n - \tilde{h}_n] \\ &= \{\xi_n(h)\}^T \Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} \{\xi_n(h)\}, \end{aligned}$$

with $\Gamma_{\mathcal{G}'} = \frac{1}{4} \left[\frac{1}{n} K_{\mathcal{G}',n} + \lambda_{\mathcal{G}'} I_n \right]^{-1}$.

The original optimization problem in h reduces to the following:

$$\hat{h} = \arg \min_{h \in \mathcal{H}_B} \{\xi_n(h)\}^T \Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} \{\xi_n(h)\} + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2.$$

By the representer theorem, the solution to this minimization problem takes the form

$$\hat{h}(a, l, w) = \sum_{j=1}^n \gamma_j K_{\mathcal{H}}\{(a, l, w), (A_j, L_j, W_j)\},$$

where $\gamma = (\gamma_1, \dots, \gamma_n)^T$ is the coefficient vector to be determined. Substituting this representation into the objective function, we obtain:

$$\frac{1}{n^2} \widetilde{Y}_n^T \Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} \widetilde{Y}_n - \frac{2}{n^2} \gamma^T K_{\mathcal{H},n} \Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} \widetilde{Y}_n + \gamma^T \left[\frac{1}{n^2} K_{\mathcal{H},n} \Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} K_{\mathcal{H},n} + \lambda_{\mathcal{H}} K_{\mathcal{H},n} \right] \gamma.$$

where $K_{\mathcal{H},n}$ is the $n \times n$ matrix with (i, j) -th entry given by $K_{\mathcal{H}} \{(A_i, L_i, W_i), (A_j, L_j, W_j)\}$. This is a convex quadratic function of γ , whose unique minimizer—found by solving the first order equation—is

$$\gamma^* = (\Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} K_{\mathcal{H},n} + n^2 \lambda_{\mathcal{H}} I_n)^{-1} \Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} \widetilde{Y}_n.$$

It remains to ensure that $\|\hat{h}\|_{\mathcal{H}} \leq B$ for some pre-specified B . If $(\gamma^*)^T K_{\mathcal{H},n} \gamma^* \leq B^2$, the condition is satisfied. Otherwise, by Exercise 4.22 of [Boyd \(2004\)](#), the optimal γ^* is given by

$$\gamma^* = \left(P + \frac{\tilde{u}}{B^2} K_{\mathcal{H},n} \right)^{-1} \tilde{q},$$

where $P = \frac{1}{n^2} K_{\mathcal{H},n} \Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} K_{\mathcal{H},n} + \lambda_{\mathcal{H}} K_{\mathcal{H},n}$, $\tilde{q} = \frac{1}{n^2} K_{\mathcal{H},n} \Gamma_{\mathcal{G}'} K_{\mathcal{G}',n} \widetilde{Y}_n$, and $\tilde{u}$ is the largest solution to the nonlinear equation:

$$\tilde{q}^T K_{\mathcal{H},n}^{-1/2} \left(K_{\mathcal{H},n}^{-1/2} P K_{\mathcal{H},n}^{-1/2} + \frac{u}{B^2} I_n \right)^{-2} K_{\mathcal{H},n}^{-1/2} \tilde{q} = B^2.$$

We now derive the closed-form expression of the estimator $\hat{g}$ for g_0 . Recall that

$$\hat{g}(a, l, z) = I_q(a, l) \cdot \tilde{g}(a, l, z),$$

where $\tilde{g}$ is the solution to the regularized minimax problem

$$\begin{aligned} \tilde{g} = \arg \min_{g \in \mathcal{G}_B} \max_{h \in \mathcal{H}'} \mathbb{E}_n \Big[& h \{q(A, L), L, W\} I \{(A, L) \in \mathcal{S}\} - I_q(A, L) h(A, L, W) g(A, L, Z) \\ & - I_q(A, L) h(A, L, W)^2 \Big] - \lambda_{\mathcal{H}'} \|h\|_{\mathcal{H}'} + \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2. \end{aligned}$$

We first show that for any $g \in \mathcal{G}$, the inner maximization satisfies

$$\max_{h \in \mathcal{H}'} \mathbb{E}_n \Big[h \{q(A, L), L, W\} I \{(A, L) \in \mathcal{S}\} - I_q(X, L) h(A, L, W) g(A, L, Z)$$

$$\begin{aligned}
& -I_q(A, L)h(A, L, W)^2] - \lambda_{\mathcal{H}'} \|h\|_{\mathcal{H}'} \\
& = \frac{1}{4} \{\zeta_n(g)\}^T \left[\frac{1}{n} K_{\mathcal{H}', n} R_n K_{\mathcal{H}', n} + \lambda_{\mathcal{H}'} K_{\mathcal{H}', n} \right]^{-1} \{\zeta_n(g)\},
\end{aligned} \tag{B.14}$$

where $K_{\mathcal{H}', n}$ is the $n \times n$ matrix with (i, j) -th entry equal to $K_{\mathcal{H}'} \{(A_i, L_i, W_i), (A_j, L_j, W_j)\}$, $\zeta_n(g) := \frac{1}{n} [\widetilde{K}_{\mathcal{H}', n}^T \tilde{d}_n - K_{\mathcal{H}', n} R_n \tilde{g}_n]$, $\tilde{d}_n$ is the vector with entries $d_i = I \{(A_i, L_i) \in \mathcal{S}\}$ for $i = 1, \dots, n$, $\tilde{g}_n = [g(A_1, L_1, Z_1), \dots, g(A_n, L_n, Z_n)]^T$, R_n is the diagonal matrix with diagonal values $r_i = I_q(A_i, L_i)$ for $i = 1, \dots, n$, and $\widetilde{K}_{\mathcal{H}', n}$ is the $n \times n$ matrix with (i, j) -th entry given by $K_{\mathcal{H}'} \{(q(A_i, L_i), L_i, W_i), (A_j, L_j, W_j)\}$.

By the generalized representer theorem, the maximizer of the optimization problem (B.14) has the form

$$h(a, l, w) = \sum_{j=1}^n \alpha_j K_{\mathcal{H}'} \{(a, l, w), (A_j, L_j, W_j)\},$$

where $\alpha = (\alpha_1, \dots, \alpha_n)^T \in \mathbb{R}^n$. The terms in the optimization problem can be expressed as follows:

$$\begin{aligned}
\mathbb{E}_n [h \{q(A, L), L, W\} I \{(A, L) \in \mathcal{S}\}] &= \frac{1}{n} \sum_{i=1}^n d_i \widetilde{K}_{\mathcal{H}', n} \alpha = \frac{1}{n} \tilde{d}_n^T \widetilde{K}_{\mathcal{H}', n} \alpha, \\
\mathbb{E}_n [I_q(A, L) h(A, L, W) g(A, L, Z)] &= \frac{1}{n} \sum_{i=1}^n g(A_i, L_i, A_i) r_i e_i^T K_{\mathcal{H}', n} \alpha = \frac{1}{n} \tilde{g}_n^T R_n K_{\mathcal{H}', n} \alpha, \\
\mathbb{E}_n [I_q(A, L) h(A, L, W)^2] &= \frac{1}{n} \sum_{i=1}^n [\alpha^T K_{\mathcal{H}', n} e_i r_i e_i^T K_{\mathcal{H}', n} \alpha] = \frac{1}{n} \alpha^T K_{\mathcal{H}', n} R_n K_{\mathcal{H}', n} \alpha, \\
\|h\|_{\mathcal{H}'}^2 &= \alpha^T K_{\mathcal{H}', n} \alpha.
\end{aligned}$$

Substituting these expressions into (B.14), we have

$$\begin{aligned}
& \max_{h \in \mathcal{H}'} \mathbb{E}_n \left[h \{q(A, L), L, W\} I \{(A, L) \in \mathcal{S}\} - I_q(A, L) h(A, L, W) g(A, L, Z) \right. \\
& \quad \left. - I_q(A, L) h(A, L, W)^2 \right] - \lambda_{\mathcal{H}'} \|h\|_{\mathcal{H}'}
\end{aligned}$$

$$= \frac{1}{n} \alpha^T [\tilde{K}_{\mathcal{H}',n}^T \tilde{d}_n - K_{\mathcal{H}',n} R_n \tilde{g}_n] - \alpha^T \left[\frac{1}{n} K_{\mathcal{H}',n} R_n K_{\mathcal{H}',n} + \lambda_{\mathcal{H}'} K_{\mathcal{H}',n} \right] \alpha.$$

The solution to this optimization problem is

$$\alpha^* = \frac{1}{2n} \left[\frac{1}{n} K_{\mathcal{H}',n} R_n K_{\mathcal{H}',n} + \lambda_{\mathcal{H}'} K_{\mathcal{H}',n} \right]^{-1} [\tilde{K}_{\mathcal{H}',n}^T \tilde{d}_n - \tilde{K}_{\mathcal{H}',n} R_n \tilde{g}_n].$$

Hence, we obtain that

$$\begin{aligned} & \max_{h \in \mathcal{H}'} \mathbb{E}_n \left[h \{q(A, L), L, W\} I \{(A, L) \in \mathcal{S}\} - I_q(A, L) h(A, L, W) g(A, L, Z) \right. \\ & \quad \left. - I_q(A, L) h(A, L, W)^2 \right] - \lambda_{\mathcal{H}'} \|h\|_{\mathcal{H}'} \\ &= \frac{1}{4n^2} [\tilde{K}_{\mathcal{H}',n}^T \tilde{d}_n - K_{\mathcal{H}',n} R_n \tilde{g}_n]^T \left[\frac{1}{n} K_{\mathcal{H}',n} R_n K_{\mathcal{H}',n} + \lambda_{\mathcal{H}'} K_{\mathcal{H}',n} \right]^{-1} [\tilde{K}_{\mathcal{H}',n}^T \tilde{d}_n - K_{\mathcal{H}',n} R_n \tilde{g}_n] \\ &= \frac{1}{4} \{\zeta_n(g)\}^T \left[\frac{1}{n} K_{\mathcal{H}',n} R_n K_{\mathcal{H}',n} + \lambda_{\mathcal{H}'} K_{\mathcal{H}',n} \right]^{-1} \{\zeta_n(g)\}. \end{aligned}$$

The original optimization problem simplifies to:

$$\tilde{g} = \arg \min_{g \in \mathcal{G}} \frac{1}{4} \{\zeta_n(g)\}^T \left[\frac{1}{n} K_{\mathcal{H}',n} R_n K_{\mathcal{H}',n} + \lambda_{\mathcal{H}'} K_{\mathcal{H}',n} \right]^{-1} \{\zeta_n(g)\} + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2.$$

By the representer theorem, the solution to this optimization problem has the form

$$\tilde{g}(a, l, z) = \sum_{j=1}^n \theta_j K_{\mathcal{G}}((a, l, z), (A_j, L_j, Z_j)),$$

where $\theta = (\theta_1, \dots, \theta_n)^T \in \mathbb{R}^n$ is the vector of coefficients to be determined. Let $\Gamma_{\mathcal{H}'} = \frac{1}{4} \left[\frac{1}{n} K_{\mathcal{H}',n} R_n + \lambda_{\mathcal{H}'} I_n \right]^{-1}$. Substituting the representation of $\tilde{g}$ into the objective function,

the objective becomes

$$\begin{aligned} & \frac{1}{4n^2} \tilde{d}_n^T \tilde{K}_{\mathcal{H}',n} \left[\frac{1}{n} K_{\mathcal{H}',n} R_n K_{\mathcal{H}',n} + \lambda_{\mathcal{H}'} K_{\mathcal{H}',n} \right]^{-1} K_{\mathcal{H}',n}^T \tilde{d}_n \\ & - \frac{2}{n^2} \theta^T K_{\mathcal{G},n} R_n \Gamma_{\mathcal{H}'} \tilde{K}_{\mathcal{H}',n}^T \tilde{d}_n + \theta^T \left[\frac{1}{n^2} K_{\mathcal{G},n} R_n \Gamma_{\mathcal{H}'} K_{\mathcal{H}',n} R_n K_{\mathcal{G},n} + \lambda_{\mathcal{G}} K_{\mathcal{G},n} \right] \theta, \end{aligned}$$

where $K_{\mathcal{G},n}$ is the $n \times n$ matrix with (i, j) -th entry equal to $K_{\mathcal{G}} \{(A_i, L_i, Z_i), (A_j, L_j, Z_j)\}$.

This unique solution to this convex quadratic function of θ is:

$$\theta^* = \left(R_n \Gamma_{\mathcal{H}'} K_{\mathcal{H}',n} R_n K_{\mathcal{G},n} + n^2 \lambda_{\mathcal{G}} I_n \right)^{-1} R_n \Gamma_{\mathcal{H}'} \tilde{K}_{\mathcal{H}',n}^T \tilde{d}_n.$$

If $(\theta^*)^T K_{\mathcal{G},n} \theta^* \leq B^2$ for some pre-specified B , it follows that $\|\tilde{g}\|_{\mathcal{G}} \leq B$. If $(\theta^*)^T K_{\mathcal{G},n} \theta^* > B^2$, the optimal θ^* ensuring $\|\tilde{g}\|_{\mathcal{G}} \leq B$ is given by

$$\theta^* = \left(P + \frac{\tilde{u}}{B^2} K_{\mathcal{G},n} \right)^{-1} \tilde{q},$$

where $P = \frac{1}{n^2} K_{\mathcal{G},n} R_n \Gamma_{\mathcal{H}'} K_{\mathcal{H}',n} R_n K_{\mathcal{G},n} + \lambda_{\mathcal{G}} K_{\mathcal{G},n}$, $\tilde{q} = \frac{1}{n^2} K_{\mathcal{G},n} R_n \Gamma_{\mathcal{H}'} \tilde{K}_{\mathcal{H}',n}^T \tilde{d}_n$, and $\tilde{u}$ is the largest solution to the nonlinear equation

$$\tilde{q}^T K_{\mathcal{G},n}^{-1/2} \left(K_{\mathcal{G},n}^{-1/2} P K_{\mathcal{G},n}^{-1/2} + \frac{u}{B^2} I_n \right)^{-2} K_{\mathcal{G},n}^{-1/2} \tilde{q} = B^2.$$

B.2.2 Estimators Accounting for Two-Phase Sampling

In the ENSEMBLE trial, the exposure—neutralizing antibody levels (A)—was measured only on a subset of study participants. For these participants, A was obtained by thawing and analyzing, at the end of follow-up, previously frozen blood samples. Subjects were selected to have A measured based on specified probabilities that depended on the outcome Y and a vector of measured baseline covariates V that included some of the variables in L .

Letting Δ_i denote the indicator that subject i is selected to have A_i measured, we then have that the observed data are n i.i.d. copies of the vector $(\Delta A, L, V, Z, W, Y)$. By design, the missingness probabilities satisfy

$$P(\Delta = 1|A, V, L, Y, Z, W) = P(\Delta = 1|V). \quad (\text{B.15})$$

Then, letting $S := P(\Delta = 1|V)^{-1}$, and with $\phi(O; h, g)$ defined in (B.10) where, recall $O := (A, L, Z, W, Y)$, we have that under the assumptions of Theorem B.1 and the equality (B.15), it holds that

$$\psi_0 := \frac{\mathbb{E} \{ \Delta S \phi(O; h^\dagger, g^\dagger) \}}{\mathbb{E} [\Delta S \cdot I \{ (A, L) \in \mathcal{S} \}]},$$

where either $h^\dagger$ is an observed outcome bridge function or $g^\dagger$ is an observed treatment bridge function, but not necessarily both. The equality, which follows by taking conditional expectation given (V, L, Z, W, Y) inside the outer expectations in the numerator and denominator, suggests the following doubly-robust cross-fitted estimator

$$\hat{\psi}_{CF}^{DR} = \frac{\sum_{k=1}^K \sum_{i: O_i \in I_k} \Delta_i S_i \phi \{ O_i; \hat{h}^{(-k)}, \hat{g}^{(-k)} \}}{\sum_{i=1}^n \Delta_i S_i I \{ (A_i, L_i) \in \mathcal{S} \}},$$

where $I_1, \dots, I_K$ is a partition of the entire sample into roughly equal size subsamples and $\hat{h}^{(-k)}$ and $\hat{g}^{(-k)}$ are estimators of the observed negative control treatment and outcome functions computed using units from the entire sample but the subsample I_k .

The next theorem extends Theorem 4.1 to the present setting with missing covariates.

Theorem B.3. *Suppose that the observed data are n i.i.d. copies of the vector $(\Delta A, L, V, Z, W, Y)$ that (B.15) holds, and that there exists $\sigma > 0$ such that $S := P(\Delta = 1|V)^{-1} < \sigma$ (a.e.). Suppose that the Assumptions of Theorem B.2 hold for some observed outcome and treatment*

bridge functions h_0 and g_0 . Then the estimator $\hat{\psi}_{CF}^{DR}$ satisfies

$$\sqrt{n} \left\{ \hat{\psi}_{CF}^{DR} - \psi_0 \right\} = \sqrt{n} \mathbb{E}_n \left[\Delta S \left\{ \frac{\phi(O; h_0, g_0)}{P \{(A, L) \in \mathcal{S}\}} - \psi_0 \frac{I \{(A, L) \in \mathcal{S}\}}{P \{(A, L) \in \mathcal{S}\}} \right\} \right] + \sqrt{n} R_n + o_p(1).$$

where $O = (A, L, Z, W, Y)$, $\phi(O; h, g)$ is defined in (B.10), and

$$R_n = \sum_{k=1}^K \frac{n_k}{n} \int \left\{ \hat{h}^{(-k)}(a, l, w) - h_0(a, l, w) \right\} \left\{ g_0(a, l, z) - \hat{g}^{(-k)}(a, l, z) \right\} dP(a, l, w, z)$$

with n_k denoting the cardinality of I_k . In particular, if $\mathbb{E} \{ \phi(O; h_0, g_0)^2 \} < \infty$ and $R_n = o_p(n^{-1/2})$, then

$$\sqrt{n} \left\{ \hat{\psi}_{CF}^{DR} - \psi_0 \right\} \xrightarrow{d} \mathcal{N}(0, \tau^2)$$

where $\tau^2 = \mathbb{E} \left\{ S \left[\frac{\phi(O; h_0, g_0)}{P \{(A, L) \in \mathcal{S}\}} - \psi_0 \frac{I \{(A, L) \in \mathcal{S}\}}{P \{(A, L) \in \mathcal{S}\}} \right]^2 \right\}$.

To construct estimators of an observed outcome bridge function we first note that for any such function h_0 it holds that

$$h_0 \in \arg \min_{h \in \mathcal{L}^2(A, L, W)} \max_{g \in \mathcal{L}^2(A, L, Z)} \mathbb{E} \left[\Delta \cdot S \left\{ g(A, L, Z) [Y - h(A, L, W)] - g(A, L, Z)^2 \right\} \right].$$

Thus, in analogy with the case in which A is always observed, we propose the estimator

$$\hat{h}^{(-k)} = \arg \min_{h \in \mathcal{H}_B} \max_{g \in \mathcal{G}'} \mathbb{E}_n^{(-k)} \left[\Delta S \left\{ g(A, L, Z) [Y - h(A, L, W)] - g(A, L, Z)^2 \right\} \right] - \lambda_{\mathcal{G}'} \|g\|_{\mathcal{G}'}^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2.$$

where, $\mathbb{E}_n^{(-k)}$ stands for the empirical mean over all sample units except those in the subsample I_k , and $\mathcal{H}$ and $\mathcal{G}'$ are the reproducing kernel Hilbert spaces described in Chapter 3 for the case that A is always observed. It can be shown that the bounds on the root-mean-squared error $\|\hat{h} - h_0\|_2$ and on the weak error $\|\mathbb{E}\{\hat{h}(A, L, W) - h_0(A, L, W) | A, L, Z\}\|_2$ established in Chapter 4 continue to hold up to constants depending on σ (as defined in Theorem B.3),

with δ_n replaced by δ_m . Here, δ_m denotes an upper bound on the critical radius—a localized measure of complexity—of certain function classes, corresponding to a sample size of m units, where m is the number of observations for which A is observed.

Likewise, for any treatment bridge function $I_q(a, l)g_0(a, l, z)$, it holds that

$$g_0 \in \arg \min_{g \in \mathcal{L}^2(A, L, Z)} \max_{h \in \mathcal{L}^2(A, L, w)} \mathbb{E} \left[\Delta \cdot S \{ h \{ q(A, L), L, W \} \cdot I \{ (A, L) \in \mathcal{S} \} \right. \\ \left. - I_q(A, L) h(A, L, W) g(A, L, Z) - I_q(A, L) h(A, L, W)^2 \} \right].$$

Hence, we propose the estimator $\hat{g} = I_q(A, L) \cdot \tilde{g}(A, L, Z)$, where

$$\tilde{g} = \arg \min_{g \in \mathcal{G}_B} \max_{h \in \mathcal{H}'} \mathbb{E}_n^{(-k)} \left[\Delta S \{ h \{ q(A, L), L, W \} I \{ (A, L) \in \mathcal{S} \} - I_q(A, L) h(A, L, W) g(A, L, Z) \right. \\ \left. - I_q(A, L) h(A, L, W)^2 \} \right] - \lambda_{\mathcal{H}'} \|h\|_{\mathcal{H}'}^2 + \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2,$$

and $\mathcal{G}$ and $\mathcal{H}'$ are the reproducing kernel Hilbert spaces described in Chapter 3.

Finally, we derive the closed-form expressions for $\hat{h}^{(-k)}$ and $\hat{g}^{(-k)}$. As in Chapter 3, we delete the superscript $(-k)$, and n stands for the cardinality of the entire sample minus the subsample I_k . Let $m = \sum_{i=1}^n \Delta_i$ denote the number of observations for which A was observed, and let $\ell_1, \dots, \ell_m$ be the indexes of those observations. Using an approach similar to the one described in Appendix B.2.1, the closed-form solution for $\hat{h}$ is given by

$$\hat{h}(a, l, w) = \sum_{j=1}^m \gamma_j^* K_{\mathcal{H}}((a, l, w), (A_{\ell_j}, L_{\ell_j}, W_{\ell_j})),$$

where

$$\gamma^* = \left(S_m \Gamma_{\mathcal{G}'} K_{\mathcal{G}', m} S_m K_{\mathcal{H}, m} + n^2 \lambda_{\mathcal{H}} I_m \right)^{-1} S_m \Gamma_{\mathcal{G}'} K_{\mathcal{G}', m} S_m \tilde{Y}_m.$$

In this expression, $K_{\mathcal{H}, m}$ and $K_{\mathcal{G}', m}$ are the $m \times m$ empirical kernel matrices evaluated at the

observations indexed by $\ell_1, \dots, \ell_m$, S_m is an $m \times m$ diagonal matrix containing the weights $s_{n_1}, \dots, s_{n_m}$ on its diagonal, $\Gamma_{\mathcal{G}'} = \frac{1}{4} \left[\frac{1}{n} S_m K_{\mathcal{G}',m} + \lambda_{\mathcal{G}'} I_m \right]^{-1}$, and $\tilde{Y}_m = (Y_{\ell_1}, \dots, Y_{\ell_m})^T$. If $(\gamma^*)^T K_{\mathcal{H},n} \gamma^* > B$ for apre-specified B , we obtain a constrained solution as outlined in Appendix B.2.1.

Similarly, the closed-form expression for $\tilde{g}$ is given by

$$\tilde{g}(a, l, z) = \sum_{j=1}^n \theta_j^* K_{\mathcal{G}}((a, l, z), (A_{\ell_j}, L_{\ell_j}, Z_{\ell_j})),$$

where the coefficients are

$$\theta^* = (S_m R_m \Gamma_{\mathcal{H}'} K_{\mathcal{H}',m} R_m S_m K_{\mathcal{G},m} + n^2 \lambda_{\mathcal{G}} I_m)^{-1} S_m R_m \Gamma_{\mathcal{H}'} \tilde{K}_{\mathcal{H}',m}^T S_m \tilde{d}_m.$$

Here, $K_{\mathcal{G},m}$ and $K_{\mathcal{H}',m}$ are the $m \times m$ empirical kernel matrices evaluated at the observations indexed by $\ell_1, \dots, \ell_m$, $\tilde{K}_{\mathcal{H}',m}$ is given by $\tilde{K}_{\mathcal{H}',m} := \left[K_{\mathcal{H}'} \left[(q(A_{\ell_i}, L_{\ell_i}), L_{\ell_i}, W_{\ell_i}), (A_{\ell_j}, L_{\ell_j}, W_{\ell_j}) \right] \right]$, R_m is an $m \times m$ diagonal matrix containing the entries $I \{A_{\ell_j} \in q(\mathcal{S}(L_{\ell_j}), L_{\ell_j})\}$ for $j = 1, \dots, m$ on its diagonal, $\Gamma_{\mathcal{H}'} = \frac{1}{4} \left[\frac{1}{n} K_{\mathcal{H}',m} S_m R_m + \lambda_{\mathcal{H}'} I_m \right]^{-1}$, and $\tilde{d}_m$ is the m -dimensional column vector with entries $d_i = I \{(A_{\ell_i}, L_{\ell_i}) \in \mathcal{S}\}$ for $i = 1, \dots, m$. If $(\theta^*)^T K_{\mathcal{G},n} \theta^* > B$, a constrained solution can be obtained as described previously. The estimator $\hat{g}$ for g_0 is $\hat{g}(a, l, z) = I \{a \in q(\mathcal{S}(l), l)\} \cdot \tilde{g}(a, l, z)$.

B.3 Additional Numerical Experiments and Implementation

Recall that $\mathcal{TN}_{[b_1, b_2]}(\mu, \sigma^2)$ denotes the probability distribution obtained from truncating a normally distributed variable $\mathcal{N}(\mu, \sigma^2)$ within the interval $[b_1, b_2]$, and that $\mathcal{TN}_c(\mu, \sigma^2)$ denotes a specific case where $b_1 = \mu - 3\sigma$ and $b_2 = \mu + 3\sigma$. Here, we present a more general data-generating process. For each $i = 1, \dots, n$, we generated variables with bounded support as follows: the observed confounder $L_i \sim \mathcal{TN}_c(0, 1)$, the unobserved confounder $U_i \sim$

$\mathcal{TN}_c(\beta_1 L_i, 1)$, the negative controls $Z_i \sim \mathcal{TN}_c(\beta_2 L_i + \beta_3 U_i, 1)$ and $W_i \sim \mathcal{TN}_c(\beta_4 L_i + \beta_5 U_i, 1)$, the exposure $A_i \sim \mathcal{TN}_{[c,d]}(\beta_6 L_i + \beta_7 U_i + \beta_8 Z_i, 1)$, and the outcome $Y_i \sim \text{Bernoulli}(p_i)$ where $p_i = [1 + \exp\{-(\mu + \beta_9 L_i + \beta_{10} U_i + \beta_{11} A_i + \beta_{12} W_i + \gamma A_i^2)\}]^{-1}$. The data-generating process described in Section 3.5 used $(\beta_1, \dots, \beta_{12}) = (0.5, 0.2, \beta_z, 0.5, \beta_w, 0.3, 1, 0, 0.5, -1, -1.5, 0)$, $c = -2$, $d = 2$, $\mu = -1$, and $\gamma = -0.75$. The data-generating process presented here additionally includes the parameter β_8 that allows the negative control treatment to have an effect on the primary treatment, and the parameter β_{12} that allows the negative control outcome affect the outcome of interest. We also consider a more general shift policy $q : [c, d - (1 - x)(\varepsilon + \delta)] \rightarrow [c + \delta, d - (1 - x)\varepsilon]$ defined as

$$q^{\delta, \varepsilon, x}(a) := a + \delta I\{c \leq a \leq d - \varepsilon - \delta\} + \frac{x \cdot \delta}{\delta + \varepsilon} (d - a) I\{d - \varepsilon - \delta < a \leq d - (1 - x)(\varepsilon + \delta)\}, \quad (\text{B.16})$$

for parameters $\delta > 0$, $\varepsilon \geq 0$, and $x \in \{0, 1\}$, satisfying $c + \delta < d - \varepsilon$. In Chapter 3, we studied the policy $q^{\delta=0.4, \varepsilon=1, x=1}$. On the other hand, policies $q^{\delta, \varepsilon, x}$ with $x = 0$ correspond to standard shift intervention policies restricted to the population $\{A \in [c, d - \delta - \varepsilon]\}$.

B.3.1 Verifying Assumption B.1–B.8

The data-generating process was constructed to reflect the structure implied by Assumptions B.1–B.5, so these assumptions are satisfied by design. Regarding Assumption B.6, expression (B.16) shows that $q^{\delta, \varepsilon, x}$ is strictly monotone and almost everywhere differentiable over $[c, d - (1 - x)(\delta + \varepsilon)]$ —that is, over $[c, d]$ when $x = 1$ and over $[c, d - \delta - \varepsilon]$ when $x = 0$. Its inverse is well-defined on $[c + \delta, d - (1 - x)\varepsilon]$ and is given by

$$(q^{\delta, \varepsilon, x})^{-1}(x) = a - \delta V(a),$$

with $V(a) = I\{c + \delta \leq a \leq d - \varepsilon\} + x \cdot \frac{d - a}{\varepsilon} I\{d - \varepsilon < a \leq d - (1 - x)\varepsilon\}$.

To verify Assumptions B.7 and B.8, we provide explicit solutions or good approximations to the bridge equations (B.2) and (B.4). Under the data-generating process, when β_{10}/β_5 is small and $\beta_{12} = 0$ (i.e., no direct effect of W on Y), a good approximation to an outcome bridge function is:

$$h(a, l, w) \approx \left[1 + \exp \left\{ \mu - \log(1 + 0.9733\varphi_3^2) + \beta_{10} a + \left(\beta_9 - \frac{\beta_4\beta_{10}}{\beta_5} \right) l + \frac{\beta_{10}}{\beta_5} w + \gamma a^2 \right\} \right]^{-1}.$$

Appendix B.6.1 provides a more general derivation that applies when β_{10}/β_5 is not small or when $\beta_{12} \neq 0$. Even though this expression is only an approximation to a solution of the outcome bridge equation (B.2), in Appendix B.3.3 we show that using this parametric model for estimation leads to excellent finite-sample performance

For the treatment bridge function, an exact solution to equation (B.4) exists when $\beta_8 = 0$ (no effect of Z on A), and is given by:

$$g(a, l, z) = \rho(a) \exp \left[\delta \left\{ a - \left(\beta_6 - \frac{\beta_2\beta_7}{\beta_3} \right) l - \frac{\beta_7}{\beta_3} z - \frac{\delta}{2} \left(1 + \frac{\beta_7^2}{\beta_3^2} \right) V(a) \right\} V(a) \right],$$

where

$$\rho(a) = \frac{\Phi(3) - \Phi(-3)}{\Phi \left\{ 3 + \frac{\delta\beta_7}{\beta_3} V(a) \right\} - \Phi \left\{ -3 + \frac{\delta\beta_7}{\beta_3} V(a) \right\}} \left(I \{ c + \delta \leq a \leq d \} + \frac{x \cdot \delta}{\varepsilon} \cdot I \{ d - \varepsilon < a \leq d \} \right),$$

and $\Phi(\cdot)$ denotes standard normal CDF.

When $\beta_8 \neq 0$, we provide an approximate solution in Appendix B.6.2.

B.3.2 Cross-Validation to Select the Bridge Functions

To implement the cross-validation procedure described in Section 3.5 of Chapter 3, we use the following empirical risk functions, defined for any functions $h \in \mathcal{L}^2(A, L, W)$ and $g \in$

$\mathcal{L}^2(A, L, Z)$ as:

$$\mathcal{R}_n^1(h) := \sup_{g \in \mathcal{G}'} \mathbb{E}_n [g(A, L, Z) \{Y - h(A, L, W)\} - g(A, L, Z)^2] - \lambda_{\mathcal{G}'} \|g\|_{\mathcal{G}'}^2,$$

and

$$\begin{aligned} \mathcal{R}_n^2(g) := & \sup_{h \in \mathcal{H}'} \mathbb{E}_n [h \{q(A, L), L, W\} \cdot I \{(A, L) \in \mathcal{S}\} \\ & - I \{A \in q(\mathcal{S}(L), L)\} \cdot h(A, L, W)g(A, L, Z) \\ & - I \{A \in q(\mathcal{S}(L), L)\} \cdot h(A, L, W)^2] - \lambda_{\mathcal{H}'} \|h\|_{\mathcal{H}'}^2. \end{aligned}$$

When the function classes $\mathcal{G}'$ and $\mathcal{H}'$ are reproducing kernel Hilbert spaces generated by Gaussian kernels, the solutions to the observed bridge equations can be shown to minimize the population version of these risk functions, making them suitable for cross-validation.

Next, we describe the cross-validation procedure when a set of candidate bridge functions are estimated using fold 2 (training set) and tested in fold 3 (testing set). Let n_3 denote the number of observations in folds 3, and let $\ell_1^{(3)}, \dots, \ell_{n_3}^{(3)}$ represent their indices. When the function classes are reproducing kernel Hilbert spaces, the empirical risk functions take the following forms:

$$\mathcal{R}_n^1(h) := \frac{1}{4} \{\xi_{n_3}(h)\}^T \left[\frac{1}{n_3} K_{\mathcal{G}', n_3}^0 + \lambda_{\mathcal{H}'} I_{n_3} \right]^{-1} K_{\mathcal{G}', n_3}^0 \{\xi_{n_3}(h)\}$$

and

$$\mathcal{R}_n^2(g) := \frac{1}{4} \{\zeta_{n_3}(g)\}^T \left[\frac{1}{n_3} K_{\mathcal{H}', n_3}^0 R_{n_3} K_{\mathcal{H}', n_3}^0 + \lambda_{\mathcal{H}'} K_{\mathcal{H}', n_3}^0 \right]^{-1} K_{\mathcal{H}', n_3}^0 \{\zeta_{n_3}(g)\}.$$

Here, $K_{\mathcal{G}', n_3}^0$ and $K_{\mathcal{H}', n_3}^0$ are $n_3 \times n_3$ matrices whose (i, j) -th entries are obtained by evaluating the kernel functions $K_{\mathcal{G}'}$ and $K_{\mathcal{H}'}$ at observations $\ell_i^{(3)}$ and $\ell_j^{(3)}$, and R_{n_3} is the diagonal matrix

with values $I\{A_{\ell_i^{(3)}} \in q(\mathcal{S}(L_{\ell_i^{(3)}}), L_{\ell_i^{(3)}})\}$ for $i = 1, \dots, n_3$ in its diagonal. The terms $\xi_{n_3}(h)$ and $\zeta_{n_3}(g)$ are defined as

$$\xi_{n_3}(h) = \frac{1}{n_3}(\tilde{Y}_{n_3} - \tilde{h}_{n_3}) \quad \text{and} \quad \zeta_{n_3}(h) = \frac{1}{n_3}\{(\tilde{K}_{\mathcal{H}', n_3}^0)^T \tilde{d}_{n_3} - K_{\mathcal{H}', n_3}^0 R_{n_3} \tilde{g}_{n_3}\},$$

where $\tilde{Y}_{n_3} = (Y_{\ell_1^{(3)}}, \dots, Y_{\ell_{n_3}^{(3)}})^T$ is the vector of observed outcomes in fold 3. The term $\tilde{h}_n = [h(A_{\ell_1^{(3)}}, L_{\ell_1^{(3)}}, W_{\ell_1^{(3)}}), \dots, h(A_{\ell_{n_3}^{(3)}}, L_{\ell_{n_3}^{(3)}}, W_{\ell_{n_3}^{(3)}})]^T$ represents the evaluations of h at the observations of fold 3, while $\tilde{g}_n = [g(A_{\ell_1^{(3)}}, L_{\ell_1^{(3)}}, Z_{\ell_1^{(3)}}), \dots, g(A_{\ell_{n_3}^{(3)}}, L_{\ell_{n_3}^{(3)}}, Z_{\ell_{n_3}^{(3)}})]^T$ represents the evaluations of g at the same observations. The matrix $\tilde{K}_{\mathcal{H}, n_3}^0$ is the $n_3 \times n_3$ is an matrix whose (i, j) -entry is given by $K_{\mathcal{H}'}^0\{(q(A_{\ell_i^{(3)}}, L_{\ell_i^{(3)}}), L_{\ell_i^{(3)}}, W_{\ell_i^{(3)}}), (A_{\ell_j^{(3)}}, L_{\ell_j^{(3)}}, W_{\ell_j^{(3)}})\}$. The bandwidth for the kernel function $K_{\mathcal{G}', n_3}^0$ is set to 1/4 times the median of the pairwise euclidean distance between the observed vectors $(A_{\ell_i}, L_{\ell_i}, Z_{\ell_i})$. The bandwidth for $K_{\mathcal{H}', n_3}^0$ is chosen analogously. The parameters $\lambda'_{\mathcal{G}}$ and $\lambda_{\mathcal{H}'}$ are set both to $\frac{\log(n_3)}{n_3}$.

B.3.3 Proximal Parametric Estimators

The data-generating mechanism suggests the following parametric model for h_0 :

$$h(a, l, w; \varphi) = \left(1 + \frac{0.9733}{2}\varphi_3^2\right) \left\{1 + \exp(-[\varphi_0 + \varphi_1 a + \varphi_2 l + \varphi_3 w + \varphi_4 a^2])\right\}^{-1}$$

and the following model for g_0 :

$$g(a, l, z; \eta) = \rho(x; \eta) \exp\{[\eta_0 a + \eta_1 l + \eta_2 z + \eta_3 V(a)] V(a)\},$$

where $V(a) = I\{c + \delta \leq a \leq d - \varepsilon\} + \frac{d-a}{\varepsilon} \cdot I\{d - \varepsilon < a \leq d\}$ and

$$\rho(a; \eta) = \frac{\Phi(3) - \Phi(-3)}{\Phi(3 - \eta_2 V(a)) - \Phi(-3 - \eta_2 V(a))} \left[I\{c + \delta \leq a \leq d\} + \frac{\delta}{\varepsilon} \cdot I\{d - \varepsilon \leq a \leq d\} \right].$$

Additionally, the integral equations (B.5) and (B.3) motivate parametric estimators $\hat{\varphi}$ for φ and $\hat{\eta}$ for η , obtained by solving the following estimating equations:

$$\begin{aligned}\mathbb{E}_n [\{Y - h(A, L, W; \varphi)\} (1, A, L, Z, A^2)^T] &= (0, 0, 0, 0, 0)^T, \\ \mathbb{E}_n [I \{(A, L) \in \mathcal{S}\} \cdot (1, q^{\delta, \varepsilon, x}(A) L, W)^T - g(A, L, Z; \eta)(1, A, L, W)^T] &= (0, 0, 0, 0)^T.\end{aligned}$$

Once estimators $\hat{\varphi}$ and $\hat{\eta}$ have been obtained, we compute the following three parametric estimators for ψ_0 :

$$\begin{aligned}\hat{\psi}_{p,OR} &= \mathbb{E}_n [h\{q^{\delta, \varepsilon, a}(A, L), W, L; \hat{\varphi}\}], \\ \hat{\psi}_{p,DQW} &= \mathbb{E}_n [Y \cdot g(A, Z, L; \hat{\eta})], \\ \hat{\psi}_{p,DR} &= \mathbb{E}_n [h\{q^{\delta, \varepsilon, a}(A, L), W, L; \hat{\varphi}\} \cdot I\{(A, L) \in \mathcal{S}\} + g(A, Z, L; \hat{\eta}) \cdot \{Y - h(A, W, L; \hat{\eta})\}].\end{aligned}$$

For each estimator, we construct 95% confidence intervals using influence-function-based standard errors that account for the estimation of the nuisance functions. Specifically, the influence functions for the target parameter under each estimator were derived using standard results from estimating equations theory, accounting for the plug-in estimation of the nuisance parameters. The asymptotic variance of each estimator is estimated via the empirical variance of the corresponding estimated influence function evaluated at the fitted nuisance models. This yields the variance estimator $\widehat{\text{var}}(\hat{\psi}) = \frac{1}{n}(\hat{\phi}_i - \hat{\psi})^2$ from which we compute Wald-type 95% confidence intervals as $\hat{\psi} \pm 1.96\sqrt{\widehat{\text{var}}(\hat{\psi})}$. These intervals are asymptotically valid under standard regularity conditions, even when the nuisance functions are estimated, provided the corresponding models are correctly specified. For $\hat{\psi}_{p,DR}$, only one of the two parametric models needs to be correctly specified to ensure consistency and valid inference.

We performed the numerical experiments for the data-generating process described in

Section 3.5. Given that the parametric approach is less computationally demanding, we considered sample sizes $n \in \{750, 3000, 12000, 48000\}$ across 16 scenarios formed by combining $\beta_3 \in \{2, 1, 0.5, 0.25\}$ and $\beta_5 \in \{-2, -1, -0.5, -0.25\}$. The values $\beta_3 = 0.25$ and $\beta_5 = -0.25$ correspond to conditional correlations $\text{Cor}(Z, U|L) = 0.235$ and $\text{Cor}(W, U|L) = -0.235$. For each configuration, we performed 500 repetitions of the data-generating and estimation procedures described here. For each nuisance function, we used a correctly specified and a misspecified parametric model. The misspecified model for h ignored the quadratic contribution of A , while the misspecified model for g replaced the term $\eta_3 V(a)$ inside by η_3 .

Figures B.1 and B.2 present boxplots and coverage probabilities of the confidence intervals for the three parametric estimators described above. Under the data-generating mechanisms considered, we find that the plug-in estimators based on correctly specified models and the double-robust estimators—provided at least one of the bridge function models is correctly specified—exhibit excellent finite-sample performance. These estimators show low bias and achieve coverage probabilities close to the nominal level. As expected, estimators relying on misspecified models are biased and have poor coverage. Furthermore, consistent with the findings in the Section 3.5 for nonparametric estimators, we observe that the challenge of estimating the bridge functions under weak conditional correlation between the negative controls and the unmeasured confounder decreases as the sample size increases




Figure B.1: Boxplots of the parametric estimators from numerical experiments with data generated under varying conditional correlations of the proxies Z and W with the unmeasured confounder U and sample sizes $n = 750, 3000, 12000$, and 48000. Dashed lines indicates the true parameter value.




Figure B.2: Empirical coverage of 95% confidence intervals from numerical experiments with data generated under varying conditional correlations of the proxies Z and W with the unmeasured confounder U and sample sizes $n = 750, 3000, 12000$, and 48000 . The dashed line indicates the nominal coverage level of 0.95.

B.3.4 Performance of the Proximal Estimator Under No Unmeasured Confounding

To assess the performance of our estimator in the absence of unmeasured confounding, we generated data under a data-generating mechanism similar to that described Section 3.5 of Chapter 3, but with β_7 set to zero. We focused solely on the scenario where $\beta_z = 2$ and $\beta_w = -2$, in which both proxies are highly correlated with the variable U . Under this configuration, Z and W act as precision variables. The target parameter for the same policy under was 0.2081. Data were generated for sample sizes $n \in \{750, 1500, 3000\}$. For each configuration, we performed 500 repetitions of the data-generating and estimation procedures. Figure B.3 displays boxplots, variance ratios, and coverage probabilities for the proximal doubly-robust cross-fitted estimator and its competitors. In this setting, both proximal and non-proximal estimators show low bias, estimated asymptotic variances closely aligned with empirical variances, and coverage probabilities at the nominal level.

B.3.5 Proximal Estimation of the Counterfactual MTP Mean in a Restricted Population

To illustrate the performance of our proximal estimator for a target parameter as described in Strategy I—outlined in Example 3.1 of Chapter 3—we conducted numerical experiments using the data-generating processes described there, but only for $\beta_z = -\beta_w \in \{2, 1, 0.5\}$. Here, however, we focus on estimating $\mathbb{E}\{Y(A + 0.4)|A \in [-2, 1.6]\}$, the counterfactual MTP mean under a shift intervention applied to the restricted population $\{A : A \in [-2, 1.6]\}$, which under all data-generating mechanisms is 0.2728. This corresponds to the policy $q^{\delta=0.4, \varepsilon=0, x=0}(a)$ defined in (B.16), for which Assumptions B.1–B.8 hold. For each setting, we performed 500 repetitions of the data-generating and estimation procedures. Figure B.4 presents boxplots, variance ratios, and coverage probabilities of the proximal doubly-robust cross-fitted estimator. Its performance is comparable to that observed for the estimand under Strategy II.




Figure B.3: Finite-sample performance of $\hat{\psi}_{CF}^{DR}$ and competitors based on numerical experiments *without* unmeasured confounding. **Panel A:** boxplots of the estimator across simulations. **Panel B:** ratio of the empirical variance to the mean estimated asymptotic variance. **Panel C:** empirical coverage of 95% confidence intervals. The dashed line indicates the true parameter value (Panel A), a variance ratio of 1 (Panel B), and nominal 95% coverage (Panel C).




Figure B.4: Finite-sample performance of $\hat{\psi}_{CF}^{DR}$ for estimating the counterfactual MTP mean in a *restricted population*, based on numerical experiments with data generated under unmeasured confounding. **Panel A:** boxplots of the estimator across simulations. **Panel B:** ratio of the empirical variance to the mean estimated asymptotic variance. **Panel C:** empirical coverage of 95% confidence intervals. The dashed line indicates the true parameter value (Panel A), a variance ratio of 1 (Panel B), and nominal 95% coverage (Panel C).

B.3.6 Performance of the Proximal Estimator Under A Data Generating Mechanism where Z affects X and W affects Y

We also conducted numerical experiments to evaluate the performance of our estimator under a data-generating mechanism where Z influences X and W affects Y . Specifically, we simulated data using the data generating process described at the beginning of Appendix B.3 with $(\beta_1, \dots, \beta_{12}) = (0.5, 0.2, \beta_z, 0.5, \beta_w, 0.3, 1, 0.3, 0.5, -1, -1.5, -0.3)$, $c = -2$, $d = 2$, $\mu = -1$, and $\gamma = -0.75$. We considered four scenarios with $\beta_z = -\beta_w \in \{3, 1.5, 1, 0.75\}$, which correspond to conditional correlations $\text{cor}(W', Z'|L) = -0.90, -0.69, -0.50$, and -0.36 , respectively. The target estimand remained the same as in the Chapter 3, namely $\mathbb{E}[Y \{q^{\delta=0.4, \varepsilon=1, x=1}(A)\}]$ which equals 0.1943 when $\beta_z = -3$, 0.2297 when $\beta_z = 1.5$, 0.2400 when $\beta_z = 1$, and 0.2466 when $\beta_z = 0.75$. Each setting is evaluated using 500 replications. We repeat the experiments for sample sizes $n \in \{750, 1500, 3000, 6000\}$.

Figure B.5 presents boxplots, variance ratios, and confidence interval coverage of the proximal doubly-robust cross-fitted estimator. Across all settings, the estimator exhibits low bias and it decreases as sample size increases. The estimated asymptotic variances closely track the empirical variances, and the confidence intervals achieve coverage near the nominal 0.95 level. As expected, estimation performance deteriorates when the conditional correlation between the negative controls and the unobserved confounder weakens, leading to increased bias and variance—consistent with the findings reported in Section 3.5 of Chapter 3 for comparable sample sizes.

B.3.7 Vignette Application to a Single Data Set

We provide R code and a mock dataset simulating data from the ENSEMBLE trial in the GitHub repository github.com/aolivasm/proxi-mtp to demonstrate the use of our `pmtpr` R function for implementing the proximal doubly-robust cross-fitted estimator.




Figure B.5: Finite-sample performance of $\hat{\psi}_{CF}^{DR}$ across varying levels of conditional correlation between the proxies Z and W and U . **Panel A:** boxplots of the estimator across simulations. **Panel B:** ratio of the empirical variance to the mean estimated asymptotic variance. **Panel C:** empirical coverage of 95% confidence intervals. The dashed line indicates the true parameter value (Panel A), a variance ratio of 1 (Panel B), and nominal 95% coverage (Panel C).

The mock dataset `sim_trial_data.csv` include the following variables:

- Outcome: `Y` (taking values 0 or 1).
- Covariates: `L1`, `L2`, `L3`.
- Biomarker: univariate biomarker `A`.
- Negative control treatment: `Z`.
- Negative control outcome: `W`.
- Sampling weights: `wt`.

The function `pmtip()` estimates the counterfactual MTP mean $\mathbb{E}[Y \{q(A)\} | A \in \mathcal{S}]$ for a policy $q(a)$ over a population of interest $\mathcal{S} \subset \text{supp}(A)$, using the proximal doubly-robust cross-fitted estimator described in Section B.1.4. It estimates the bridge functions using reproducing kernel Hilbert spaces generated by Gaussian kernels as outlined in sections B.2.1, B.2.2, and B.3.2. The function supports both binary and continuous outcomes, but does not handle time-to-event (censored) outcomes. Additionally, `pmtip()` only uses complete cases. For data arising from a two-phase sampling design, providing NA values for the treatment variable for observations outside the reduced sample and supplying the corresponding sampling weight is sufficient.

The `pmtip()` function requires the following inputs:

- `data`: A `data.frame` with the dataset to analyze.
- `trt`: String for the biomarker (treatment) variable name.
- `outcome`: String for outcome variable name.
- `covariates`: String or vector of strings for covariate(s) name(s).
- `nct`: String or vector of strings for negative control treatment variable(s) name(s).
- `nco`: String or vector of strings for negative control outcome variable(s) name(s).

- **policy**: A list of policy functions. Functions should depend only on the treatment variable.

The following inputs are optional. If not provided, default values will be used:

- **weights**: String for weights variable name.
- **ind_S**: A numeric vector of 0s and 1s indicating membership in the target population $\mathcal{S}$. Default is a vector of 1s (i.e., all units are in $\mathcal{S}$).
- **theta**: Scaling factor c_0 for the regularization parameter $c_0 \cdot \frac{\log(n)}{n}$ in the risk function used during cross-validation to select the optimal bridge functions. Default is 1.
- **K_folds**: Integer indicating the number of folds used for cross-fitting. Must be at least 3. Default is 3.
- **lm_H_list**, **lm_G_list**: Lists of scaling factors c_1 for the regularization parameters $c_1 \left(\frac{\log(n)}{n} \right)^{1/2}$ in the outer minimization step of the outcome and treatment bridge function estimation, respectively. Defaults are `10^seq(-5, -1, by = 1)`.
- **lm_Gh_list**, **lm_Hg_list**: Lists of scaling factors c_2 for the regularization parameters $c_2 \frac{\log(n)}{n}$ in the inner maximization step of the outcome and treatment bridge function estimation, respectively. Defaults are `10^seq(-1, 2, by = 1)`.
- **bw0_H**, **bw0_G**: Optional initial bandwidths for the function classes $\mathcal{H}$ and $\mathcal{H}'$, and $\mathcal{G}$ and $\mathcal{G}'$, respectively. Default is NULL. If not provided, the initial bandwidths are set to the median (or weighted median, if weights are provided) of the pairwise Euclidean distances between the corresponding observed vectors.
- **bw_int_scale_list**, **bw_ext_scale_list**: Lists of scaling factors for the bandwidths of the Gaussian kernels used in the internal and external steps of the bridge function estimation procedure. Defaults are 1/4 and `c(1/4, 1/2, 1, 2, 4)`, respectively.

- `bw.int.fixed.scale`: Scale factor used for the internal bandwidth of the Gaussian kernel used in the risk function for cross-validation. Default is `1/4`.
- `inf.fun`: Logical flag indicating whether to return the influence function evaluated at the sample observations. Default is `FALSE`.
- `show.prog`: Logical flag indicating whether to display progress messages. Default is `FALSE`.
- `control.folds`: Logical flag indicating whether fold assignment in cross-fitting should be stratified based on the distribution of the treatment variable. Default is `FALSE`.

The function outputs a `data.frame` with the proximal doubly-robust cross-fitted estimator for each specified policy and also returns a `proximal.mtp` object with detailed information about the estimation procedure. In addition, the `summary.proximal.mtp()` function reports, for each specified policy, the proportion of observations with treatment value within the image of the policy, the proportion of observations whose treatment values fall within the image of the policy, along with standard errors and 95% confidence intervals for the estimators.

Input to conduct analysis on the mock dataset

We illustrate our estimator for the following two policies:

```
policy_q1 <- function(x) {
  (x+0.4)*(x+0.4<=3.5-1)+(0.4*3.5+1*x)/(0.4+1)*(x+0.4>3.5-1)
}

policy_q2 <- function(x) {
  (x+0.8)*(x+0.8<=3.5-1)+(0.8*3.5+1*x)/(0.8+1)*(x+0.8>3.5-1)
}

obj.pmtmp = pmtmp(data = sim_trial_data,
```

```

trt = "A",
outcome = "Y",
covariates = c("L1", "L2", "L3"),
nct = "Z",
nco = "W",
weights = "wt",
policy = list(policy_q1, policy_q2),
control.folds = TRUE)

```

Output using the summary.proximal.mtp() function

```
summary(obj.pmt)
```

Summary of Proximal MTP Estimation

```

-----
Number of observations: 1000
Number in target population: 1000 (100%)
-Proportion in the image of Policy_q^1: 38%
-Proportion in the image of Policy_q^2: 27%

```

```

-----
Proximal Estimators and 95% Confidence Intervals:
-----

```

Estimator: Proximal doubly-robust cross-fitted

	Estimate	Std.Error	CI.Lower	CI.Upper
Policy_q^1	0.0139	7e-04	0.0126	0.0152
Policy_q^2	0.0101	7e-04	0.0088	0.0114

B.4 Proofs of Main Results

Proof of Theorem 1. Proof of part i). Take $(a, l, z) \in \text{supp}(A, L, Z)$ such that $a \in \mathcal{S}(l) \cup q(\mathcal{S}(l), l)$. Hence, if $h_0(a, l, w)$ solves the integral equation (B.2), we have

$$\begin{aligned}
\int h_0(a, l, w) p_{W|A,L,Z}(w|a, l, z) dw &= \int h_0(a, l, w) \left\{ \int p_{(W,U)|A,L,Z}(w, u|a, l, z) du \right\} dw \\
&= \int h_0(a, l, w) \int p_{W|A,L,U}(w|a, l, u) p_{U|A,L,Z}(u|a, l, z) dudw \quad \text{by Assumption B.5} \\
&= \int \left\{ \int h_0(a, l, w) p_{W|A,L,U}(w|a, l, u) dw \right\} p_{U|A,L,Z}(u|a, l, z) du \\
&= \int \mathbb{E}(Y|A = a, L = l, U = u) p_{U|A,L,Z}(u|a, l, z) du \quad h_0 \text{ solves (B.2)} \\
&= \int \mathbb{E}(Y|A = a, L = l, U = u, Z = z) p_{U|A,L,Z}(u|a, l, z) du \quad \text{by Assumption B.4} \\
&= \mathbb{E}(Y|A = a, L = l, Z = z).
\end{aligned}$$

Thus, $h_0(a, l, w)$ solves the integral equation (B.5).

Proof of part ii). Take $(a, l, w) \in \text{supp}(A, L, W)$ satisfying $x \in \mathcal{S}(l) \cup q(\mathcal{S}(l), l)$, and such that, if $a \in q(\mathcal{S}(l), l)$, the map $x' \mapsto q^{-1}(a', l)$ is differentiable at $a' = a$. Hence, if $g_0(a, l, z)$ solves the integral equation (B.4), we have

$$\begin{aligned}
\int g_0(a, l, z) p_{Z|A,L,W}(z|a, l, w) dz &= \int g_0(a, l, z) \left\{ \int p_{(Z,U)|A,L,W}(z, u|a, l, w) du \right\} dz \\
&= \int g_0(a, l, z) \int p_{Z|A,L,U}(z|a, l, u) p_{U|A,L,W}(u|a, l, w) dudz \quad \text{by Assumption B.5} \\
&= \int \left\{ \int g_0(a, l, z) p_{Z|A,L,U}(z|a, l, u) dz \right\} p_{U|A,L,W}(u|a, l, w) du \\
&= \int \alpha_0(a, l, u) p_{U|A,L,W}(u|a, l, w) du \quad g_0 \text{ solves (B.4)} \\
&= \int I\{a \in q(\mathcal{S}(l), l)\} \frac{dq^{-1}(a, l)}{da} \cdot \frac{p_{A|L,U}\{q^{-1}(a, l)|l, u\}}{p_{A|L,U}(a|l, u)} p_{U|A,L,W}(u|a, l, w) du \\
&\stackrel{(i)}{=} I\{a \in q(\mathcal{S}(l), l)\} \frac{dq^{-1}(a, l)}{da} \int \frac{p_{U|A,L,W}(u|a, l, w)}{p_{A|L,U,W}(a|l, u, w)} p_{A|L,U,W}\{q^{-1}(a, l)|l, u, w\} du
\end{aligned}$$

$$\begin{aligned}
&= I\{a \in q(\mathcal{S}(l), l)\} \frac{dq^{-1}(a, l)}{da} \int \frac{p_{U|L,W}(u|l, w)}{p_{A|L,W}(a|l, w)} p_{A|L,U,W}\{q^{-1}(a, l) | l, u, w\} du \\
&= I\{a \in q(\mathcal{S}(l), l)\} \frac{dq^{-1}(a, l)}{da} \frac{1}{p_{A|L,W}(a|l, w)} \int p_{(A,U)|L,W}\{q^{-1}(a, l), u|l, w\} du \\
&= I\{a \in q(\mathcal{S}(l), l)\} \frac{dq^{-1}(a, l)}{da} \frac{p_{A|L,W}\{q^{-1}(a, l) | l, w\}}{p_{A|L,W}(a|l, w)} \\
&= \alpha_0(a, l, w),
\end{aligned}$$

where the equality in (i) follows from Assumption B.5. Thus, $g_0(a, l, z)$ solves equation (B.3). $\square$

Theorem B.1. Case 1. Assumption B.7 holds and equation (B.3) has at least one solution.

Let h_0 and $g^\dagger$ be solutions to equations (B.2) and (B.3), respectively. Then,

$$\begin{aligned}
\psi_0 &= \mathbb{E}[Y\{q(A, L)\} | (A, L) \in \mathcal{S}] \\
&= \frac{\mathbb{E}[Y\{q(A, L)\} \cdot I\{(A, L) \in \mathcal{S}\}]}{P\{(A, L) \in \mathcal{S}\}} \\
&= \frac{\mathbb{E}\{\mathbb{E}[Y\{q(A, L)\} | A, L, U] \cdot I\{(A, L) \in \mathcal{S}\}\}}{P\{(A, L) \in \mathcal{S}\}} \\
&= \frac{\mathbb{E}\{\mathbb{E}[h_0\{q(A, L), L, W\} | A, L, U] \cdot I\{(A, L) \in \mathcal{S}\}\}}{P\{(A, L) \in \mathcal{S}\}} \quad \text{by Lemma B.3} \\
&= \frac{\mathbb{E}[h_0\{q(A, L), L, W\} \cdot I\{(A, L) \in \mathcal{S}\}]}{P\{(A, L) \in \mathcal{S}\}}.
\end{aligned}$$

By Theorem 3.1, h_0 also solves equation (B.5), then, by Lemma B.2,

$$\frac{\mathbb{E}\{Y g^\dagger(A, L, Z)\}}{P\{(A, L) \in \mathcal{S}\}} = \frac{\mathbb{E}[h_0\{q(A, L), L, W\} \cdot I\{(A, L) \in \mathcal{S}\}]}{P\{(A, L) \in \mathcal{S}\}} = \psi_0,$$

which is the treatment bridge representation.

Now, let $h^\dagger$ be any solution to equation (B.5). By Lemma B.2,

$$\frac{\mathbb{E} [h^\dagger \{q(A, L), L, W\} \cdot I \{(A, L) \in \mathcal{S}\}]}{P \{(A, L) \in \mathcal{S}\}} = \frac{\mathbb{E} [Y g^\dagger(A, L, Z)]}{P \{(A, L) \in \mathcal{S}\}} = \psi_0,$$

which is the outcome bridge representation.

Lastly, for any $h \in \mathcal{L}^2(A, L, W)$ and any $g \in \mathcal{L}^2(A, L, Z)$, by Lemma B.2,

$$\mathbb{E} \{\phi(O; h, g^\dagger)\} = \mathbb{E} \{\phi(O; h^\dagger, g^\dagger)\} = \mathbb{E} [h^\dagger \{q(A, L), L, W\} \cdot I \{(A, L) \in \mathcal{S}\}],$$

and

$$\mathbb{E} \{\phi(O; h^\dagger, g)\} = \mathbb{E} \{\phi(O; h^\dagger, g^\dagger)\} = \mathbb{E} \{Y g^\dagger(A, L, Z)\}.$$

Hence

$$\frac{\mathbb{E} \{\phi(O; h, g^\dagger)\}}{P \{(A, L) \in \mathcal{S}\}} = \psi_0 = \frac{\mathbb{E} \{\phi(O; h^\dagger, g)\}}{P \{(A, L) \in \mathcal{S}\}},$$

which is the double robust representation.

Case 2. Assumption B.8 holds and equation (B.5) has at least one solution. Let g_0 and $h^\dagger$ be solutions to equations (B.4) and (B.5), respectively. Then,

$$\begin{aligned} \psi_0 &= \mathbb{E}[Y \{q(A, L)\} | (A, L) \in \mathcal{S}] \\ &= \frac{\mathbb{E}\{\mathbb{E}[Y \{q(A, L)\} | A, L, U] \cdot I \{(A, L) \in \mathcal{S}\}\}}{P \{(A, L) \in \mathcal{S}\}} \\ &= \iiint \frac{I \{(a, l) \in \mathcal{S}\}}{P \{(A, L) \in \mathcal{S}\}} \mathbb{E}\{Y | A = q(a, l), L = l, U = u\} p_{A,L,U}(a, l, u) da dl du \\ &= \iint \left[\int \frac{I \{(a, l) \in \mathcal{S}\}}{P \{(A, L) \in \mathcal{S}\}} \mathbb{E}\{Y | A = q(a, l), L = l, U = u\} p_{A|L,U}(a|l, u) da \right] p_{L,U}(l, u) dl du \\ &\stackrel{(i)}{=} \iiint \frac{I \{\tilde{a} \in q(\mathcal{S}(l), l)\}}{P \{(A, L) \in \mathcal{S}\}} \mathbb{E}(Y | A = \tilde{a}, L = l, U = u) \frac{dq^{-1}(\tilde{a}, l)}{d\tilde{a}} p_{A|L,U}\{q^{-1}(\tilde{a}, l) | l, u\} p_{L,U}(l, u) d\tilde{a} dl du \\ &= \frac{1}{P \{(A, L) \in \mathcal{S}\}} \iiint \mathbb{E}[Y | A = \tilde{a}, L = l, U = u] \cdot \alpha_0(\tilde{a}, l, u) \cdot p_{A,L,U}(\tilde{a}, l, u) d\tilde{a} dl du \end{aligned}$$

$$\begin{aligned}
&\stackrel{(ii)}{=} \frac{1}{P\{(A, L) \in \mathcal{S}\}} \mathbb{E} [\mathbb{E}(Y|A, L, U) \cdot \mathbb{E}\{g_0(A, L, Z)|A, L, U\}] && g_0 \text{ solves (B.4)} \\
&= \frac{\mathbb{E}\{Y g_0(A, L, Z)\}}{P\{(A, L) \in \mathcal{S}\}},
\end{aligned}$$

where the equality in (i) follows from Assumption B.6, which justifies applying the change of variable $\tilde{a} = q(a, l)$ in a for each l , and from the fact that $(a, l) \in \mathcal{S}$ if and only if $q(a, l) \in q(\mathcal{S}(l), l)$. For the equality in (ii), note that $\tilde{a} \in q(\mathcal{S}(l), l)$ implies $\tilde{a} = q(a, l)$ for some $a \in \mathcal{S}(l)$, and by the definitions of $\mathcal{S}$ and $\mathcal{S}(l)$, $\tilde{a} \in \text{supp}(A|L = l)$. Then, for any $u \in \text{supp}(U|L)$, Assumption B.3 implies $\tilde{a} \in \text{supp}(A|L = l, U = u)$, and consequently $(\tilde{a}, l, u) \in \text{supp}(A, L, U)$. Therefore, $(\tilde{a}, l, u)$ lies in the domain where g_0 solves equation (B.4). By Theorem 3.1, g_0 also solves equation (B.3). Thus, by Lemma B.2,

$$\frac{\mathbb{E} [h^\dagger \{q(A, L), L, W\} \cdot I\{(A, L) \in \mathcal{S}\}]}{P\{(A, L) \in \mathcal{S}\}} = \frac{\mathbb{E}\{Y g_0(A, L, Z)\}}{P\{(A, L) \in \mathcal{S}\}} = \psi_0,$$

which is the outcome bridge representation. From here, the other two representations are derived analogously to Case 1. $\square$

Proof of Theorem B.2. Let $\hat{p} := \frac{1}{n} \sum_{i=1}^n I\{(A_i, L_i) \in \mathcal{S}\}$ and

$$U_n := \frac{1}{\sqrt{n}} \sum_{k=1}^K \sum_{i: O_i \in I_k} \left[\phi\{O_i; \hat{h}^{(-k)}, \hat{g}^{(-k)}\} - \mathbb{E}\{\phi(O; h_0, g_0)\} \right].$$

Noting that $\psi_0 = p_0^{-1} \mathbb{E}\{\phi(O; h_0, g_0)\}$, where $p_0 := P\{(A, L) \in \mathcal{S}\}$, we have

$$\begin{aligned}
\sqrt{n} (\hat{\psi}_{CF}^{DR} - \psi_0) &:= \sqrt{n} \left[\frac{1}{\hat{p}} \cdot \frac{1}{n} \sum_{k=1}^K \sum_{i: O_i \in I_k} \phi\{O; \hat{h}^{(-k)}, \hat{g}^{(-k)}\} - \psi_0 \right] \\
&= \frac{1}{p_0} U_n + \left(\frac{1}{\hat{p}} - \frac{1}{p_0} \right) U_n + \sqrt{n} \left(\frac{1}{\hat{p}} - \frac{1}{p_0} \right) \mathbb{E}\{\phi(O; h_0, g_0)\}. \tag{B.17}
\end{aligned}$$

By the Delta Method,

$$\sqrt{n}(\hat{p}^{-1} - p_0^{-1}) = -p_0^{-2} \mathbb{G}_n I \{(A, L) \in \mathcal{S}\}, \quad (\text{B.18})$$

where for any $t(O)$, $\mathbb{G}_n(t) = \sqrt{n}[n^{-1} \sum_i^n t(O_i) - \mathbb{E}\{t(O)\}]$. Likewise, letting n_k denote the cardinality of I_k , $k = 1, \dots, K$, we let $\mathbb{E}_{n_k}(t) := n_k^{-1} \sum_{i: O_i \in I_k} t(O_i)$ denote the empirical mean of $t(O)$ over the sample I_k , and $\mathbb{G}_{n_k}(t) := \sqrt{n_k} \{\mathbb{E}_{n_k}(t) - \int t(o) dP(o)\}$. Then

$$\begin{aligned} U_n &= \sum_{k=1}^K \sqrt{\frac{n_k}{n}} \mathbb{G}_{n_k} \left\{ \phi(\cdot; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(\cdot; h_0, g_0) \right\} + \sum_{k=1}^K \sqrt{\frac{n_k}{n}} \mathbb{G}_{n_k} \left\{ \phi(\cdot; h_0, g_0) \right\} \\ &\quad + \sqrt{n} \sum_{k=1}^K \frac{n_k}{n} \int \left\{ \phi(o; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(o; h_0, g_0) \right\} dP(o). \end{aligned}$$

Now, let $I^{(-k)} := \{O_i : 1 \leq i \leq n, O_i \notin I_k\}$ and let m_k denote the cardinality of $I^{(-k)}$. For any $\epsilon > 0$, we have

$$\begin{aligned} D_{m_k} &:= P \left[\left| \mathbb{G}_{n_k} \left\{ \phi(\cdot; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(\cdot; h_0, g_0) \right\} \right| \geq \epsilon \middle| I^{(-k)} \right] \\ &\leq \mathbb{E} \left[\left\{ \mathbb{G}_{n_k} \left\{ \phi(\cdot; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(\cdot; h_0, g_0) \right\} \right\}^2 \middle| I^{(-k)} \right] / \epsilon^2 \\ &\leq \mathbb{E} \left[\left\{ \phi(O; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(O; h_0, g_0) \right\}^2 \middle| I^{(-k)} \right] / \epsilon^2 \\ &= \int \left\{ \phi(o; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(o; h_0, g_0) \right\}^2 dP(o) \\ &= o_P(1), \end{aligned}$$

where the first inequality is by Markov's inequality, the second is because given $I^{(-k)}$, $\mathbb{G}_{n_k}^{(k)} \left\{ \phi(\cdot; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(\cdot; h_0, g_0) \right\}$ is a sum of n_k independent and identically distributed random variables, the second to last equality is by independence of $I^{(-k)}$ and I_k and the last equality follows from Lemma B.4. We then conclude that $\mathbb{E}(D_{m_k})$ converges to 0 as $m_k \rightarrow \infty$

because D_{m_k} , $m_k = 1, 2, \dots$ is a sequence of positive random variables bounded above by 1 that converges in probability to 0. But

$$\begin{aligned}\mathbb{E}(D_{m_k}) &= \mathbb{E} \left\{ P \left[\left| \mathbb{G}_{n_k} \left\{ \phi(\cdot; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(\cdot; h_0, g_0) \right\} \right| \geq \epsilon \middle| I^{(-k)} \right] \right\} \\ &= P \left[\left| \mathbb{G}_{n_k} \left\{ \phi(\cdot; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(\cdot; h_0, g_0) \right\} \right| \geq \epsilon \right].\end{aligned}$$

Then, assuming as we do, that $n_k/n \rightarrow 1/K$ as $n \rightarrow \infty$ we have that $m_k \rightarrow \infty$ as $n \rightarrow \infty$, and therefore $\sum_{k=1}^K \sqrt{\frac{n_k}{n}} \mathbb{G}_{n_k} \left\{ \phi(\cdot; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(\cdot; h_0, g_0) \right\} = o_p(1)$ as $n \rightarrow \infty$.

Next, by Lemma B.2, for each k , it holds that

$$\begin{aligned}& \int \left\{ \phi(o; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(o; h_0, g_0) \right\} dP(o) \\ &= \int \left\{ \hat{h}^{(-k)}(a, l, w) - h_0(a, l, w) \right\} \left\{ g_0(a, l, z) - \hat{g}^{(-k)}(a, l, z) \right\} dP(a, l, w, z).\end{aligned}$$

Thus, $\sum_{k=1}^K \frac{n_k}{n} \int \left\{ \phi(o; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(o; h_0, g_0) \right\} dP(o) = p_0 R_n$ where R_n is defined in the statement of the theorem. From (C.1) we thus conclude that

$$U_n = \mathbb{G}_n \left\{ \phi(\cdot; h_0, g_0) \right\} + \sqrt{n} p_0 R_n + o_p(1) \quad (\text{B.19})$$

By the Cauchy-Schwarz inequality and the norm consistency assumption on $h^{(-k)}$ and $g^{(-k)}$, we have

$$|R_n| \leq p_0^{-1} \sum_{k=1}^K \frac{n_k}{n} \|\hat{h}^{(-k)} - h_0\|_2 \|\hat{g}^{(-k)} - g_0\|_2 = o_p(1).$$

Because $\mathbb{G}_n \left\{ \phi(\cdot; h_0, g_0) \right\} = O_p(1)$ and (B.18) implies that $\hat{p}^{-1} - p_0^{-1} = O_p(n^{-1/2})$, it follows that

$$(\hat{p}^{-1} - p_0^{-1}) U_n = o_p(1). \quad (\text{B.20})$$

Replacing the equalities (B.18), (B.19), and (B.20) into (B.17), we arrive at

$$\begin{aligned}
\sqrt{n} (\psi_{CF}^{DR} - \psi_0) &= p_0^{-1} \mathbb{G}_n \{ \phi(\cdot; h_0, g_0) \} + \sqrt{n} R_n \\
&\quad - p_0^{-2} \mathbb{G}_n I \{ (A, L) \in \mathcal{S} \} \mathbb{E} \{ \phi(O; h_0, g_0) \} + o_p(1) \\
&= p_0^{-1} \mathbb{G}_n \left\{ \phi(\cdot; h_0, g_0) - I \{ (A, L) \in \mathcal{S} \} \psi_0 \right\} + \sqrt{n} R_n + o_p(1).
\end{aligned}$$

This concludes the proof of the Theorem. $\square$

B.5 Supporting Lemmas

In this section, we present Lemmas that are used in the proofs of the results provided in B.1.

Lemma B.1 below implies that $\alpha_0(A, L, W)$ is the Riesz representer of the linear functional given by $h \mapsto \mathbb{E} [h \{ q(A, L), L, W \} \cdot I \{ (A, L) \in \mathcal{S} \}]$.

Lemma B.1. *Under Assumption B.6, for any $h \in \mathcal{L}^2(A, L, W)$, it holds that*

$$\mathbb{E} [h \{ q(A, L), L, W \} \cdot I \{ (A, L) \in \mathcal{S} \}] = E \{ h(A, L, W) \alpha_0(A, L, W) \}.$$

Proof. For any $h \in \mathcal{L}^2(A, L, W)$, we have

$$\begin{aligned}
&\iiint h(a, l, w) \alpha_0(a, l, w) p_{A,L,W}(a, l, w) da dl dw \\
&= \iiint h(a, l, w) I \{ a \in q(\mathcal{S}(l), l) \} \frac{dq^{-1}(a, l)}{da} \frac{p_{A|L,W} \{ q^{-1}(a, l) | l, w \}}{p_{A|L,W}(a | l, w)} p_{A,L,W}(a, l, w) da dl dw \\
&= \iint \left[h(a, l, w) I \{ a \in q(\mathcal{S}(l), l) \} \frac{dq^{-1}(a, l)}{da} p_{A|L,W} \{ q^{-1}(a, l) | l, w \} da \right] p_{L,W}(l, w) dl dw \\
&\stackrel{(i)}{=} \iiint h \{ q(\tilde{a}, l), l, w \} I \{ (\tilde{a}, l) \in \mathcal{S} \} \frac{dq^{-1} \{ q(\tilde{a}, l), l \}}{d\tilde{a}} \cdot \frac{dq(\tilde{a}, l)}{d\tilde{a}} p_{A|L,W}(\tilde{a} | l, w) p_{L,W}(l, w) d\tilde{a} dl dw \\
&\stackrel{(ii)}{=} \iiint h \{ q(\tilde{a}, l), l, w \} I \{ (\tilde{a}, l) \in \mathcal{S} \} p_{A,L,W}(\tilde{a}, l, w) d\tilde{a} dl dw,
\end{aligned}$$

where the equality in (i) follows from Assumption B.6, which allows us to apply the change of variable $\tilde{a} = q^{-1}(a, l)$ in a for each $l \in \text{supp}(L)$, and from the fact that $a \in q(\mathcal{S}(l), l)$ if and only if $(q^{-1}(a, l), l) \in \mathcal{S}$. For the equality in (ii), we used that for each $l \in \text{supp}(L)$,

$$1 = \frac{d\tilde{a}}{d\tilde{a}} = \frac{d[q^{-1}\{q(\tilde{a}, l), l\}]}{d\tilde{a}} = \frac{dq^{-1}\{q(\tilde{a}, l), l\}}{da} \cdot \frac{dq(\tilde{a}, l)}{d\tilde{a}}.$$

This concludes the proof. $\square$

Lemma B.2. *Let $h^\dagger$ and $g^\dagger$ be observed outcome and treatment bridge functions respectively. Under Assumption B.6, for any $h \in \mathcal{L}^2(A, L, W)$ and $g \in \mathcal{L}^2(A, L, Z)$, it holds that*

$$\mathbb{E}\{\phi(O; h, g) - \phi(O; h^\dagger, g^\dagger)\} = \mathbb{E}\left[\{(h - h^\dagger)(A, L, W)\} \{(g^\dagger - g)(A, L, Z)\}\right].$$

Furthermore,

$$\mathbb{E}[h^\dagger\{q(A, L), L, W\} \cdot I\{(A, L) \in \mathcal{S}\}] = \mathbb{E}\{\phi(O; h^\dagger, g^\dagger)\} = \mathbb{E}\{Y g^\dagger(A, L, Z)\}.$$

Proof. Let $h^\dagger$ and $g^\dagger$ be observed outcome and treatment bridge functions respectively; and take any $h \in \mathcal{L}^2(A, L, W)$ and $g \in \mathcal{L}^2(A, L, Z)$. Then,

$$\begin{aligned} & \mathbb{E}\{\phi(O; h, g) - \phi(O; h^\dagger, g^\dagger)\} \\ &= \mathbb{E}[(h - h^\dagger)\{q(A, L), L, W\} \cdot I\{(A, L) \in \mathcal{S}\}] + \mathbb{E}[g(A, L, Z)\{Y - h(A, L, W)\}] \\ & \quad - \mathbb{E}\left[g^\dagger(A, L, Z) \underbrace{\mathbb{E}\{Y - h^\dagger(A, L, W)|A, L, Z\}}_{=0 \text{ since } h^\dagger \text{ is an outcome br fcn}}\right] \\ &= \mathbb{E}\{(h - h^\dagger)(A, L, W)\alpha_0(A, L, W)\} + \mathbb{E}\{g(A, L, Z) \cdot \mathbb{E}(Y|A, L, Z)\} \quad \text{by Lemma B.1} \\ & \quad - \mathbb{E}\{g(A, L, Z) \cdot h(A, L, W)\} \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E} \left[(h - h^\dagger)(A, L, W) \underbrace{\mathbb{E} \{g^\dagger(A, L, Z) | A, L, W\}}_{g^\dagger \text{ is a trx br fcn}} \right] + \mathbb{E} \left[g(A, L, Z) \underbrace{\mathbb{E} \{h^\dagger(A, L, W) | A, L, Z\}}_{h^\dagger \text{ is an outcome br fcn}} \right] \\
&\quad - \mathbb{E} [g(A, L, Z) \cdot h(A, L, W)] \\
&= \mathbb{E} \left[\{(h - h^\dagger)(A, L, W)\} \{(g^\dagger - g)(A, L, Z)\} \right].
\end{aligned}$$

Taking $h \equiv h^\dagger$ and $g \equiv 0$ gives

$$\mathbb{E} [h^\dagger \{q(A, L), L, W\} \cdot I \{(A, L) \in \mathcal{S}\}] = \mathbb{E} \{\phi(O; h^\dagger, 0)\} = \mathbb{E} \{\phi(O; h^\dagger, g^\dagger)\},$$

and taking $h \equiv 0$ and $g \equiv g^\dagger$ gives

$$\mathbb{E} \{Y g^\dagger(A, L, Z)\} = \mathbb{E} \{\phi(O; 0, g^\dagger)\} = \mathbb{E} \{\phi(O; h^\dagger, g^\dagger)\}.$$

This concludes the proof. □

Lemma B.3. *Under Assumptions B.1–B.3, for any $(a, l, u) \in \text{supp}(A, L, U)$ such that $a \in \mathcal{S}(l)$, it holds that*

$$\mathbb{E} [Y \{q(A, L)\} | A = a, L = l, U = u] = \mathbb{E} \{Y | A = q(a, l), L = l, U = u\}.$$

Moreover, if Assumption B.5 holds, then for any latent outcome bridge function h_0 it holds that

$$\mathbb{E} [Y \{q(A, L)\} | A = a, L = l, U = u] = \mathbb{E} [h_0 \{q(A, L), L, W\} | A = a, L = l, U = u].$$

Proof. Take $(a, l, u) \in \text{supp}(A, L, U)$ such that $a \in \mathcal{S}(l)$. By the definitions of $\mathcal{S}$ and $\mathcal{S}(l)$,

$q(a, l) \in \text{supp}(A|L = l)$. Then, by Assumption B.3, $q(a, l) \in \text{supp}(A|L = l, U = u)$, which implies that $(q(a, l), l, u) \in \text{supp}(A, L, U)$. Then,

$$\begin{aligned} \mathbb{E}[Y \{q(A, L)\} | A = a, L = l, U = u] &= \mathbb{E}[Y \{q(a, l)\} | A = a, L = l, U = u] \\ &= \mathbb{E}[Y \{q(a, l)\} | A = q(a, l), L = l, U = u] \\ &= \mathbb{E}\{Y | A = q(a, l), L = l, U = u\}, \end{aligned}$$

where the second equality follows from Assumption B.2 and the last one from Assumption B.1.

Now, let h_0 be a solution to equation (B.2). Since $(q(a, l), l, u) \in \text{supp}(A, L, U)$ and $q(a, l) \in q(\mathcal{S}(l), l)$, $(q(a, l), l, u)$ lies in the domain where h_0 solves equation (B.2). Then,

$$\begin{aligned} \mathbb{E}[Y | A = q(a, l), L = l, U = u] &= \int h_0\{q(a, l), l, w\} p_{W|A, L, U}\{w | q(a, l), l, u\} dw \\ &= \int h_0\{q(a, l), l, w\} p_{W|A, L, U}(w | a, l, u) dw \quad \text{by Assumption B.5} \\ &= \mathbb{E}[h_0\{q(A, L), L, W\} | A = a, L = l, U = u]. \end{aligned}$$

This concludes the proof. □

Lemma B.4. *Under the assumptions of Theorem B.2, for each k , it holds that*

$$\|\phi(O; \hat{h}^{(-k)}, \hat{g}^{(-k)}) - \phi(O; h_0, g_0)\|_2 = o_p(1).$$

Proof. From the definition of ϕ in (B.10), we have

$$\phi(O; \hat{h}, \hat{g}) - \phi(O; h_0, g_0) = (\hat{h} - h_0) \{q(A, L), L, W\} \cdot I\{(A, L) \in \mathcal{S}\} + Y(\hat{g} - g_0)(A, L, Z)$$

$$- \hat{h}(A, L, W) \hat{g}(A, L, Z) + h_0(A, L, W) g_0(A, L, Z),$$

and by the triangle inequality

$$\begin{aligned} \sqrt{\mathbb{E} \left[\left\{ \phi(O; \hat{h}, \hat{g}) - \phi(O; h_0, g_0) \right\}^2 \right]} &\leq \sqrt{\mathbb{E} \left[\left\{ (\hat{h} - h_0) \{q(A, L), L, W\} \cdot I \{(A, L) \in \mathcal{S}\} \right\}^2 \right]} \\ &\quad + \sqrt{\mathbb{E} \left[\{Y (\hat{g} - g_0)(A, L, Z)\}^2 \right]} \\ &\quad + \sqrt{\mathbb{E} \left[\left\{ \hat{h}(A, L, W) \hat{g}(A, L, Z) - h_0(A, L, W) g_0(A, L, Z) \right\}^2 \right]} \end{aligned}$$

Since $|\alpha_0(A, L, W)| \leq B$ it follows that

$$\sqrt{\mathbb{E} \left[\left\{ (\hat{h} - h_0) \{q(A, L), L, W\} \cdot I \{(A, L) \in \mathcal{S}\} \right\}^2 \right]} \leq \sqrt{B} \|\hat{h} - h_0\|_2 = O(\|\hat{h} - h_0\|_2).$$

This result is proven in Appendix C.1 of Chapter 4. Also, since $|Y| \leq B$, it follows that

$$\sqrt{\mathbb{E} \left[\{Y (\hat{g} - g_0)(A, L, Z)\}^2 \right]} \leq B \|\hat{g} - g_0\|_2 = O(\|\hat{g} - g_0\|_2).$$

Furthermore, the term

$$\sqrt{\mathbb{E} \left[\left\{ \hat{h}(A, L, W) \hat{g}(A, L, Z) - h_0(A, L, W) g_0(A, L, Z) \right\}^2 \right]} = \|\hat{h} \cdot \hat{g} - h_0 \cdot g_0\|_2$$

is upper bounded by

$$\begin{aligned} &\min \left\{ \|\hat{h}(\hat{g} - g_0)\|_2 + \|g_0(\hat{h} - h_0)\|_2, \|\hat{g}(\hat{h} - h_0)\|_2 + \|h_0(\hat{g} - g_0)\|_2 \right\} \\ &\leq \min \left\{ \|\hat{h}\|_\infty \|\hat{g} - g_0\|_2 + \|g_0\|_\infty \|\hat{h} - h_0\|_2, \|\hat{g}\|_\infty \|\hat{h} - h_0\|_2 + \|h_0\|_\infty \|\hat{g} - g_0\|_2 \right\} \\ &= O \left(\|\hat{h} - h_0\|_2 + \|\hat{g} - g_0\|_2 \right) \end{aligned}$$

where the last equality follows because by assumption $\|\hat{h}\|_\infty + \|g_0\|_\infty \leq B$ or $\|\hat{g}\|_\infty + \|h_0\|_\infty \leq B$.

Thus, we conclude that

$$\sqrt{\mathbb{E} \left[\left\{ \phi(O; \hat{h}, \hat{g}) - \phi(O; h_0, g_0) \right\}^2 \right]} = O \left(\|\hat{h} - h_0\|_2 + \|\hat{g} - g_0\|_2 \right)$$

which is $o_p(1)$ since $\|\hat{h} - h_0\|_2^2 = o_p(1)$, $\|\hat{g} - g_0\|_2^2 = o_p(1)$.

□

B.6 Existence of Bridge Functions for our Data Generating Mechanism

B.6.1 Derivation of an Outcome Bridge Function

Under our data-generating mechanism, the outcome bridge equation (B.5) becomes

$$\mathbb{E}(Y|A, L, U) = \left[1 + \exp \left\{ -(\mu + \beta_9 L + \beta_{10} U + \beta_{11} A + \beta_{12} W + \gamma A^2) \right\} \right]^{-1}.$$

This structural form motivates the following parametric model for h :

$$h(A, L, W; \varphi) = \kappa(\varphi) \left[1 + \exp \left\{ -(\varphi_0 + \varphi_1 A + \varphi_2 L + \varphi_3 W + \varphi_4 A^2) \right\} \right]^{-1}.$$

First, consider the expression

$$\varphi_0 + \varphi_1 A + \varphi_2 L + \varphi_3 W + \varphi_4 A^2 = \varphi_0 + \varphi_3 \widetilde{W} + \varphi_1 A + (\varphi_2 + \varphi_3 \beta_4) L + \varphi_3 \beta_5 U + \varphi_4 A^2,$$

where $\widetilde{W} = W - \beta_4 L - \beta_5 U$ is a random variable with distribution $\mathcal{TN}_c(0, 1)$. We treat $\widetilde{W}$ as a mean-zero perturbation and use a second-order Taylor expansion to approximate the effect of W on the expected value of $h(A, L, W; \varphi)$ given (A, L, U) . Specifically, we consider the

function $f(v, t) = v \{1 + \exp(-t)\}^{-1}$, which second-order Taylor approximation at (v_0, t_0) is given by:

$$\begin{aligned}
f(v, t) &\approx f(v_0, t_0) + (v - v_0) \frac{\partial f(v_0, t_0)}{\partial v} + (t - t_0) \frac{\partial f(v_0, t_0)}{\partial t} + \frac{(t - t_0)^2}{2} \frac{\partial^2 f(v_0, t_0)}{\partial t^2} \\
&\quad + (v - v_0)(t - t_0) \frac{\partial^2 f(v_0, t_0)}{\partial v \partial t} \\
&= f(v_0, t_0) + (v - v_0) \frac{f(v_0, t_0)}{v_0} + (t - t_0) f(v_0, t_0) \left\{ 1 - \frac{f(v_0, t_0)}{v_0} \right\} \\
&\quad + \frac{(t - t_0)^2}{2} f(v_0, t_0) \left\{ 1 - \frac{f(v_0, t_0)}{v_0} \right\} \left\{ 1 - \frac{2f(v_0, t_0)}{v_0} \right\} \\
&\quad + (v - v_0)(t - t_0) \frac{f(v_0, t_0)}{v_0} \left\{ 1 - \frac{f(v_0, t_0)}{v_0} \right\}
\end{aligned}$$

Applying the Taylor approximation with $v_0 = 1$, $t_0 = \varphi_0 + \varphi_1 A + (\varphi_2 + \varphi_3 \beta_4) L + \varphi_3 \beta_5 U + \varphi_4 A^2$, $v = \kappa(\varphi)$, and $t = t_0 + \varphi_3 \widetilde{W}$, and taking expectations with respect to (A, L, U) yields the following expression:

$$\begin{aligned}
&\mathbb{E} \{h(A, L, W; \varphi) | A, L, U\} \\
&= \mathbb{E} \{f(1, t) | A, L, U\} \\
&\approx f(1, t_0) + \{\kappa(\varphi) - 1\} f(1, t_0) + \frac{\varphi_3^2}{2} f(1, t_0) \{1 - f(1, t_0)\} \{1 - 2f(1, t_0)\} \mathbb{E}[\widetilde{W}^2] \\
&\approx \left\{ \kappa(\varphi) + \frac{\varphi_3^2}{2} \mathbb{E}[\widetilde{W}^2] \right\} f(1, t_0) \\
&= \left\{ \kappa(\varphi) + \frac{\varphi_3^2}{2} \mathbb{E}[\widetilde{W}^2] \right\} \left[1 + \exp \left\{ -(\varphi_0 + \varphi_1 A + (\varphi_2 + \varphi_3 \beta_4) L + \varphi_3 \beta_5 U + \varphi_4 A^2) \right\} \right]^{-1}.
\end{aligned}$$

Using a similar analysis, an approximation of $\mathbb{E}(Y | X, L, U)$ is given by

$$\left(1 + \frac{\beta_{12}^2}{2} \mathbb{E}[\widetilde{W}^2] \right) \left[1 + \exp \left\{ -(\mu + \beta_{11} A + (\beta_9 + \beta_{12} \beta_4) L + (\beta_{10} + \beta_{12} \beta_5) U + \gamma A^2) \right\} \right]^{-1}.$$

This suggests the following condition must be approximately satisfied to match both sides

of the integral equation (3.1):

$$\begin{aligned} & \left(1 + \frac{\beta_{12}^2}{2} \mathbb{E}[\widetilde{W}^2]\right) [1 + \exp \{-(\mu + \beta_{11}A + (\beta_9 + \beta_{12}\beta_4)L + (\beta_{10} + \beta_{12}\beta_5)U + \gamma_4 A^2)\}]^{-1} \\ & \approx \left\{ \kappa(\varphi) + \frac{\varphi_3^2}{2} \mathbb{E}[\widetilde{W}^2] \right\} [1 + \exp \{-(\varphi_0 + \varphi_1 A + (\varphi_2 + \varphi_3\beta_4)L + \varphi_3\beta_5 U + \varphi_4 A^2)\}]^{-1}, \end{aligned}$$

where $\mathbb{E}[\widetilde{W}^2] \approx 0.9733369$ is a constant.

This approximation suggests $(\varphi_1, \varphi_2, \varphi_3, \varphi_4) \approx (\beta_{11}, \beta_9 - \frac{\beta_4\beta_{10}}{\beta_5}, \beta_{12} + \frac{\beta_{10}}{\beta_5}, \gamma)$ and $\kappa(\varphi) \approx 1 + (\beta_{12}^2 - \varphi_3^2)\mathbb{E}[\widetilde{W}^2]/2$. When $\beta_{12} = 0$, i.e., when W does not affect Y , i.e., the previous approximation provides a parametric form for $\kappa(\varphi)$ provided $\mathbb{E}[\widetilde{W}^2]$ is known, an assumption that we make. Unfortunately, this choice can be problematic as $\kappa(\varphi)$ may become negative when β_5 is small—i.e., when W is a weak proxy for U —leading to unstable estimates. To address this issue, we instead choose $\kappa(\varphi) = 1 + \frac{\varphi_3^2}{2} \mathbb{E}[\widetilde{W}^2]$. Under this selection, when β_{10}/β_5 is small, setting

$$\varphi_0 \approx \mu + \log \left(1 + \frac{\beta_{12}^2}{2} \mathbb{E}[\widetilde{W}^2]\right) - \log \left\{1 + \left(\beta_{12} + \frac{\beta_{10}}{\beta_5}\right)^2 \mathbb{E}[\widetilde{W}^2]\right\}$$

ensures that the proposed function h approximately solves the integral equation (3.1). In the general case, using the same model for $\kappa(\varphi)$, i.e., $\kappa(\varphi) = 1 + \frac{\varphi_3^2}{2} \mathbb{E}[\widetilde{W}^2]$ with $\mathbb{E}[\widetilde{W}^2] \approx 0.9733369$, an approximation of φ can be found by solving a non-linear system of equations. We used such equations for the parametric estimation φ described in Section B.3.3. Our numerical experiments described in Appendix B.3.3 showed that this choice of $\kappa(\varphi)$ yields good finite-sample performance for estimating ψ_0 when $\beta_{12} = 0$. We also observed good performance under a more challenging setting with $\beta_9 = 0.3$ and $\beta_{12} = -0.3$ (see Figures B.6 and B.7).

B.6.2 Derivation of a Treatment Bridge Function

Under our data-generating mechanism and policy (B.16), when $\beta_8 = 0$, the expression for $\alpha_0(a, l, u)$ becomes

$$\begin{aligned}\alpha_0(a, l, u) &= I\{a \in q[\mathcal{S}(l)]\} \frac{dq^{-1}(a, l)}{da} \frac{p_{A|L,U}(q^{-1}(a, l)|l, u)}{p_{A|L,U}(a|l, u)} \\ &= \rho_0(a) \exp \left[-\frac{1}{2} \{q^{-1}(a) - a\} \{q^{-1}(a) + a - 2(\beta_6 l + \beta_7 u)\} \right] \\ &= \rho_0(a) \exp \left[\frac{\delta V(a)}{2} \{2a - \delta V(a) - 2(\beta_6 l + \beta_7 u)\} \right] \\ &= \rho_0(a) \exp \left[\delta V(a) \left\{ a - \beta_6 l - \beta_7 u - \frac{\delta}{2} V(a) \right\} \right]\end{aligned}$$

where $\rho_0(a) = I\{c + \delta \leq a \leq d\} + \frac{\delta}{\varepsilon} I\{d - \varepsilon \leq a \leq d\}$ and, recall, $V(a) = I\{c + \delta \leq a < d - \varepsilon\} + \frac{d-a}{\varepsilon} \cdot I\{d - \varepsilon \leq a \leq d\}$. Since the latent treatment bridge equation (B.3) has $\alpha_0(a, l, u)$ on its left-hand side, the expression above suggests the following parametric model for g :

$$g(a, l, z; \eta) = \rho(a; \eta) \exp [\{\eta_0 a + \eta_1 l + \eta_2 z + \eta_3 V(a)\} V(a)].$$

Noting that the moment generating function of $Z|A, L, U$ is

$$\mathbb{E} \{\exp(tZ)|A, L, U\} = \frac{\Phi(3-t) - \Phi(-3-t)}{\Phi(3) - \Phi(-3)} \exp \left([\beta_2 L + \beta_3 U] t + \frac{t^2}{2} \right),$$

where t could be a function of (a, l, u) , it follows that

$$\begin{aligned}\mathbb{E} \{g(A, L, Z; \eta)|A, L, U\} &= \rho(a; \eta) \exp [\{\eta_0 X + \eta_1 L + \eta_3 V(X)\} V(X)] \cdot \mathbb{E} \{\eta_2 V(X)Z|X, L, U\} \\ &= \rho_1(a; \eta) \exp \left[\left\{ \eta_0 A + (\eta_1 + \eta_2 \beta_2) L + \eta_2 \beta_3 U + \left(\eta_3 + \frac{\eta_2^2}{2} \right) V(A) \right\} V(A) \right],\end{aligned}$$

where $\rho_1(a; \eta) = \frac{\Phi(3-\eta_2 V(a)) - \Phi(-3-\eta_2 V(a))}{\Phi(3) - \Phi(-3)} \rho(a; \eta)$.

Hence, taking

$$\rho(a; \eta) = \frac{\Phi(3) - \Phi(-3)}{\Phi\{3 - \eta_2 V(a)\} - \Phi\{-3 - \eta_2 V(a)\}} \cdot \rho_0(a)$$

and

$$\eta = \left[\delta, -\delta \left(\beta_6 - \frac{\beta_2 \beta_7}{\beta_3} \right), -\delta \frac{\beta_7}{\beta_3}, -\frac{\delta^2}{2} \left(1 + \frac{\beta_7^2}{\beta_3^2} \right) \right]^T,$$

ensures that the proposed function $\hat{g}(a, l, u; \eta)$ solves the integral equation (3.3).

When $\beta_8 \neq 0$, using the identity $p_{A|L,U}(a|l, u) = \int p_{A|L,U,Z}(a|l, u, z) p_{Z|L,U}(z|l, u) dz$, it can be shown that

$$\alpha_0(a, l, u) = \tilde{\rho}_0(a) \exp \left[\frac{\delta}{1 + \beta_8^2} V(a) \left\{ a - (\beta_6 + \beta_2 \beta_8) l - (\beta_7 + \beta_3 \beta_8) u - \frac{\delta}{2} V(a) \right\} \right],$$

where

$$\tilde{\rho}_0(x) = \rho_0(a) \cdot \frac{\theta(a - \delta V(a), l, u)}{\theta(a, l, u)},$$

and

$$\theta(a, l, u) := \int_{-3+\beta_1+\beta_3 u}^{3+\beta_2 l+\beta_3 u} \frac{\exp \left(-\frac{1+\beta_8^2}{2} \left[z - \frac{1}{1+\beta_8^2} \{ \beta_8 a + (\beta_2 - \beta_6 \beta_8) l + (\beta_3 - \beta_7 \beta_8) u \}^2 \right] \right)}{\Phi(2 - \beta_6 l - \beta_7 u - \beta_8 z) - \Phi(-2 - \beta_6 l - \beta_7 u - \beta_8 z)} dz.$$

Similarly, it can be shown that

$$\begin{aligned} \mathbb{E}[g(A, L, Z; \eta) | A, L, U] &= \\ &= \rho(a; \eta) \exp [\{ \eta_0 A + \eta_1 L + \eta_3 V(A) \} V(A)] \cdot \mathbb{E}[\eta_2 V(A) Z | A, L, U] \\ &= \rho_1(a; \eta) \exp \left(\left[\left(\eta_0 + \frac{\eta_2 \beta_8}{1 + \beta_8^2} \right) A + \left\{ \eta_1 + \frac{\eta_2 (\beta_2 - \beta_6 \beta_8)}{1 + \beta_8^2} \right\} L \right. \right. \end{aligned}$$

$$+ \frac{\eta_2(\beta_3 - \beta_7\beta_8)}{1 + \beta_8^2} U + \left\{ \eta_3 + \frac{\eta_2^2}{2(1 + \beta_8^2)} \right\} V(A) \Big] V(A) \Big),$$

where

$$\tilde{\rho}_1(a; \eta) := \rho(a; \eta) \cdot \frac{\theta\left(a + \frac{\eta_2 V(a)}{\beta_8}, l, u\right)}{\theta(a, l, u)}.$$

Now define, for fixed (a, l, u) , $r(\beta_8) := \frac{\theta(a - \delta V(a), l, u)}{\theta(a + \frac{\eta_2 V(a)}{\beta_8}, l, u)}$. Since $r'(0) \approx 0$ under the parameters considered in our data-generating process, a first-order Taylor approximation gives

$$r(\beta_8) \approx r(0) = \frac{\Phi(3) - \Phi(-3)}{\Phi\{3 - \eta_2 V(a)\} - \Phi\{-3 - \eta_2 V(a)\}}.$$

Hence, taking

$$\eta = \left[\frac{\delta\beta_3}{\beta_3 - \beta_7\beta_8}, -\frac{\delta(\beta_3\beta_6 - \beta_2\beta_7)}{\beta_3 - \beta_7\beta_8}, -\frac{\delta(\beta_7 + \beta_3\beta_8)}{\beta_3 - \beta_7\beta_8}, -\frac{\delta^2}{(\beta_3 - \beta_7\beta_8)^2} \left\{ \frac{\beta_3^2 + \beta_7^2}{2} \right\} \right],$$

and

$$\rho(a; \eta) = \frac{\Phi(3) - \Phi(-3)}{\Phi\{3 - \eta_2 V(a)\} - \Phi\{-3 - \eta_2 V(a)\}} \cdot \rho_0(a),$$

yields a good approximation to the true bridge function when β_8 is small. In the general case, using the same model for $\rho(a; \eta)$, an approximation of η can be found by solving the non-linear system of equations described in Section B.3.3. Our numerical experiments, using parametric models for estimating the bridge functions under a more challenging data-generating mechanism with $\beta_8 = 0.3$ and $\beta_{12} = -0.3$ demonstrated that the proposed parametric model for g yields good finite-sample performance for estimating ψ_0 .




Figure B.6: Boxplots of the *correctly specified* parametric estimators from numerical experiments under a more challenging data generating mechanism, where Z directly affects A and W directly affects Y . Results are shown across varying conditional correlations of the proxies Z and W with the unmeasured confounder U , and sample sizes $n = 750, 3000, 12000$, and 48000 . Dashed lines indicate the true parameter values.




Figure B.7: Empirical coverage of 95% confidence intervals of *correctly specified* parametric estimators from numerical experiments under a more challenging data generated-generating mechanism, where Z directly affects A and W directly affects Y . Results are shown across varying conditional correlations of the proxies Z and W with the unmeasured confounder U , and sample sizes $n = 750, 3000, 12000$, and 48000 . Dashed lines indicate the true parameter values. The dashed line indicates the nominal coverage level of 0.9.

Appendix C

SUPPLEMENTARY MATERIAL OF CHAPTER 4

C.1 Proofs of Mathematical Statements

Proof of Proposition 4.1. If $\{g_j = I_0 g_j^\dagger : g_j^\dagger \in \mathcal{G}, j \in \mathbb{J}\}$ is an orthonormal basis of $\overline{I_0 \mathcal{G}}$, then

$$\|\mathcal{P}_{\overline{I_0 \mathcal{G}}} \circ \mathcal{T} I_0(h_0^\dagger - h)\|_2^2 = \sum_{j \in \mathbb{J}} \langle I_0 \mathcal{T}(h_0^\dagger - h), I_0 g_j^\dagger \rangle_2^2 = \sum_{j \in \mathbb{J}} \{\mathbb{E}[\rho(Z) g_j^\dagger(Z) - I_0(X) r(O) h(W) g_j^\dagger(Z)]\}^2 = \sum_{j \in \mathbb{J}} \{\mathbb{E}[m(O, g_j^\dagger) - I_0(X) r(O) h(W) g_j^\dagger(Z)]\}^2 = Q_{\mathbb{J}}(h) \quad \square$$

Proof of Proposition 4.2. To show part 1, write

$$\begin{aligned} M_P(h) &\stackrel{(i)}{=} 4c^2 \max_{g \in \mathcal{G}} \mathbb{E} \{ \rho(Z) g(Z) - I_0(X) r(O) h(W) g(Z) - c^2 I_0(X) g(Z)^2 \} \\ &\stackrel{(ii)}{=} 4c^2 \max_{g \in \mathcal{G}} \mathbb{E} \left\{ I_0(X) g(Z) \mathcal{T}(h_0^\dagger - h)(Z) - c^2 I_0(X) g(Z)^2 \right\} \\ &= 4c^2 \max_{g \in \mathcal{G}} \mathbb{E} \left[I_0(X) g(Z) \mathcal{T}(h_0^\dagger - h)(Z) - c^2 I_0(X) g(Z)^2 - \frac{1}{4c^2} I_0(X) \{ \mathcal{T}(h_0^\dagger - h)(Z) \}^2 \right] \\ &\quad + \mathbb{E} \left[I_0(X) \{ \mathcal{T}(h_0^\dagger - h)(Z) \}^2 \right] \\ &= \max_{g \in \mathcal{G}} -\mathbb{E} \left[I_0(X) \{ 2c^2 g(Z) - \mathcal{T}(h_0^\dagger - h)(Z) \}^2 \right] + \mathbb{E} \left[I_0(X) \{ \mathcal{T}(h_0^\dagger - h)(Z) \}^2 \right] \\ &\stackrel{(iii)}{=} \mathbb{E} \left[I_0(X) \{ \mathcal{T}(h_0^\dagger - h)(Z) \}^2 \right] \end{aligned}$$

where the equality (i) holds because ρ is the Riesz representer of the map $g \mapsto \mathbb{E}[m(O; g)]$, the equality (ii) is because $I_0 h_0^\dagger$ is a solution to equation (4.2), and the equality (iii) is because by assumption, $(2c^2)^{-1} \cdot \mathcal{T}(h_0^\dagger - h) \in \mathcal{G}$. Note that if $\mathcal{G}$ was unrestricted, part 1 would be, on expectation, the Fenchel conjugate representation of the squared function.

This concludes the proof of part 1. To show part 2 write

$$\begin{aligned}
& \max_{g \in \mathcal{G}} -\mathbb{E} \left[I_0(X) \{ 2c^2 g(Z) - \mathcal{T}(h_0^\dagger - h)(Z) \}^2 \right] + \mathbb{E} \left[I_0(X) \{ \mathcal{T}(h_0^\dagger - h)(Z) \}^2 \right] \\
& \stackrel{(i)}{=} -\min_{g \in \mathcal{G}} \mathbb{E} \left[\{ I_0(X) g(Z) - I_0(X) \mathcal{T}(h_0^\dagger - h)(Z) \}^2 \right] + \mathbb{E} \left[I_0(X) \{ \mathcal{T}(h_0^\dagger - h)(Z) \}^2 \right] \\
& \stackrel{(ii)}{=} -\| I_0 \mathcal{T}(h_0^\dagger - h) - \mathcal{P}_{\overline{I_0 \mathcal{G}}} \circ \mathcal{T} I_0(h_0^\dagger - h) \|_2^2 + \| I_0 \mathcal{T}(h_0^\dagger - h) \|_2^2 \\
& \stackrel{(iii)}{=} \| \mathcal{P}_{\overline{I_0 \mathcal{G}}} \circ \mathcal{T} I_0(h_0^\dagger - h) \|_2^2
\end{aligned}$$

where the equality (i) holds because $\mathcal{G}$ is a linear space, the equality (ii) holds by the definition of the projection operator and the assumption that $\mathcal{P}_{\overline{I_0 \mathcal{G}}} \circ \mathcal{T} I_0(h_0^\dagger - h) \in I_0 \mathcal{G}$, and the equality (iii) holds by Pythagorean Theorem. Hence, by Proposition 4.1, we have $M_P(h) = \| \mathcal{P}_{\overline{I_0 \mathcal{G}}} \circ \mathcal{T} I_0(h_0^\dagger - h) \|_2^2 = Q_{\mathbb{J}}(h)$. This concludes the proof of part 2. $\square$

Theorem 4.3 under Assumption 4.4. Recall that $h'_{0,KRAS}$ is the solution to equation (4.17) with minimal $\mathcal{L}^2(P_W)$ -norm and that $h'_{\lambda_{\mathcal{H}},KRAS} = \arg \min_{h' \in \mathcal{N}(T_{\mathcal{H}})^\perp} \{ \| \rho - \widetilde{\mathcal{T}} h' \|_2^2 + 4\lambda_{\mathcal{H}} \| h' \|_2^2 \}$ (see Section 4.5). Now, Lemma C.4 establishes that $h'_{0,KRAS} \in \mathcal{N}(T_{\mathcal{H}})^\perp$ and $h_{0,KRAS}^\dagger = T_{\mathcal{H}}^{1/2} h'_{0,KRAS}$. On the other hand, we know that $h'_{\lambda_{\mathcal{H}},KRAS} \in \mathcal{N}(T_{\mathcal{H}})^\perp$ and $h_{\lambda_{\mathcal{H}},KRAS}^\dagger = T_{\mathcal{H}}^{1/2} h'_{\lambda_{\mathcal{H}},KRAS}$. It then follows from Lemma C.1 (d) that $\| h_{\lambda_{\mathcal{H}},KRAS}^\dagger - h_{0,KRAS}^\dagger \|_{\mathcal{H}} = \| h'_{\lambda_{\mathcal{H}},KRAS} - h'_{0,KRAS} \|_2$ and $\| \widetilde{\mathcal{T}}(h'_{\lambda_{\mathcal{H}},KRAS} - h'_{0,KRAS}) \|_2 = \| I_0^{op} \circ \mathcal{T}(h_{\lambda_{\mathcal{H}},KRAS}^\dagger - h_{0,KRAS}^\dagger) \|_2$. Hence, under Assumption 4.4, Theorem 4.2 with $\mathcal{A} = \widetilde{\mathcal{T}}$ ensures that there exist constants c_1^* and c_2^* such that

$$\begin{aligned}
& \| h_{\lambda_{\mathcal{H}},KRAS}^\dagger - h_{0,KRAS}^\dagger \|_{\mathcal{H}}^2 = \| h'_{\lambda_{\mathcal{H}},KRAS} - h'_{0,KRAS} \|_2^2 \leq c_1^* \lambda_{\mathcal{H}}^{\min\{\beta, 2\}} \quad \text{and} \\
& \| I_0^{op} \circ \mathcal{T}(h_{\lambda_{\mathcal{H}},KRAS}^\dagger - h_{0,KRAS}^\dagger) \|_2^2 = \| \widetilde{\mathcal{T}}(h'_{\lambda_{\mathcal{H}},KRAS} - h'_{0,KRAS}) \|_2^2 \leq c_2^* \lambda_{\mathcal{H}}^{\min\{\beta+1, 2\}}.
\end{aligned}$$

$\square$

Theorem 4.4. We adapted the proof of Theorem 4 of [Bennett et al. \(2023b\)](#) to our estimator. To simplify notation, we omit the subscript $KRAS$ from the functions $h_{0,KRAS}^\dagger$, $h_{\lambda_{\mathcal{H}},KRAS}^\dagger$, and $\tilde{h}_{KRAS}$ defined in (4.13), (4.15), and (4.12). Define

$$L(\tau) := \frac{1}{4c^2} \|I_0 \mathcal{T}\{h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger - \tau(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\}\|_2^2 + \lambda_{\mathcal{H}} \|h_{\lambda_{\mathcal{H}}}^\dagger + \tau(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\|_{\mathcal{H}}^2.$$

$L(\tau)$ attains its unique global minimum at $\tau = 0$ because $h_{\lambda_{\mathcal{H}}}^\dagger$ is the unique minimizer in $\mathcal{H}$ of the criterion in (4.15). Then, $L'(0) = 0$ and since $L(\tau)$ is quadratic, a Taylor's expansion yields

$$L(1) - L(0) = L'(0) \cdot 1 + \frac{1}{2} L''(0) \cdot 1^2 = \frac{1}{2} L''(0) = \frac{1}{4c^2} \|I_0 \mathcal{T}(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + \lambda_{\mathcal{H}} \|\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger\|_{\mathcal{H}}^2.$$

On the other hand, $L(0) = \frac{1}{4c^2} \|I_0 \mathcal{T}(h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + \lambda_{\mathcal{H}} \|h_{\lambda_{\mathcal{H}}}^\dagger\|_{\mathcal{H}}^2$ and $L(1) = \frac{1}{4c^2} \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 + \lambda_{\mathcal{H}} \|\tilde{h}\|_{\mathcal{H}}^2$. Thus, we arrive at

$$\begin{aligned} & \frac{1}{4c^2} \|I_0 \mathcal{T}(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + \lambda_{\mathcal{H}} \|\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger\|_{\mathcal{H}}^2 \\ &= \frac{1}{4c^2} \left\{ \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 - \|I_0 \mathcal{T}(h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 \right\} + \lambda_{\mathcal{H}} \left(\|\tilde{h}\|_{\mathcal{H}}^2 - \|h_{\lambda_{\mathcal{H}}}^\dagger\|_{\mathcal{H}}^2 \right). \end{aligned} \quad (\text{C.1})$$

To control the right-hand side of (C.1), we invoke Lemma C.10 which establishes that with probability at least $\kappa(\delta_n)$ the following event happens

$$\begin{aligned} & \frac{1}{4c^2} \left\{ \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 - \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 \right\} + \lambda_{\mathcal{H}} \left(\|\tilde{h}\|_{\mathcal{H}}^2 - \|h\|_{\mathcal{H}}^2 \right) \\ & \leq \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 + c_0 \delta_n \|I_0 \mathcal{T}(\tilde{h} - h)\|_2 + \lambda_{\mathcal{G}} B_1^2 + c'_0 \delta_n^2 \end{aligned}$$

where $c_0 = 9 \cdot \max\{\|s\|_\infty, \|r\|_\infty \cdot k_1 B, 2c^2 b\} / (c^2 b)$ and $c'_0 = c^2 \cdot c_0(2_3 c_0 + \max\{2, 2c_0/9\})$.

Hence, identity (C.1) implies that, with probability at least $\kappa(\delta_n)$, the following holds

$$\begin{aligned} \frac{1}{4c^2} \|I_0 \mathcal{T}(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + \lambda_{\mathcal{H}} \|\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger\|_{\mathcal{H}}^2 \\ \leq \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + c_0 \delta_n \|I_0 \mathcal{T}(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2 + \lambda_{\mathcal{G}} B_1^2 + c_0' \delta_n^2. \end{aligned}$$

By the AM-GM inequality, $c_0 \delta_n \|I_0 \mathcal{T}(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2 \leq 2c^2 \cdot c_0^2 \delta_n^2 + \frac{1}{8c^2} \|I_0 \mathcal{T}(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2$, which combined with the right-hand side of the last display yields that

$$\lambda_{\mathcal{H}} \|\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger\|_{\mathcal{H}}^2 + \frac{1}{8c^2} \|I_0 \mathcal{T}(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 \leq \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + \lambda_{\mathcal{G}} B_1^2 + c_0'' \delta_n^2, \quad (\text{C.2})$$

with $c_0'' = c_0' + 2c^2 \cdot c_0^2$, happens with probability at least $\kappa(\delta_n)$, which, in turn, implies that $\lambda_{\mathcal{H}} \|\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger\|_{\mathcal{H}}^2 \leq \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + \lambda_{\mathcal{G}} B_1^2 + c_0'' \delta_n^2$ also holds with probability at least $\kappa(\delta_n)$.

Applying the inequality $\|h_1 + h_2\|_{\mathcal{H}}^2 \leq 2\|h_1\|_{\mathcal{H}}^2 + 2\|h_2\|_{\mathcal{H}}^2$ we then obtain that the event

$$\begin{aligned} \|\tilde{h} - h_0^\dagger\|_{\mathcal{H}}^2 &\leq 2\|\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger\|_{\mathcal{H}}^2 + 2\|h_{\lambda_{\mathcal{H}}}^\dagger - h_0^\dagger\|_{\mathcal{H}}^2 \\ &\leq \frac{1}{c^2 \cdot \lambda_{\mathcal{H}}} \|I_0 \mathcal{T}(h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + \frac{2\lambda_{\mathcal{G}} B_1^2}{\lambda_{\mathcal{H}}} + \frac{2c_0'' \delta_n^2}{\lambda_{\mathcal{H}}} + 2\|h_{\lambda_{\mathcal{H}}}^\dagger - h_0^\dagger\|_{\mathcal{H}}^2, \end{aligned}$$

happens with probability at least $\kappa(\delta_n)$. Because $\tilde{h} - h_0^\dagger \in \mathcal{H}$, it follows that $\|\tilde{h} - h_0^\dagger\|_2^2 \leq k_1^2 \|\tilde{h} - h_0^\dagger\|_{\mathcal{H}}^2$. Hence, invoking Theorem 4.3 to bound the terms $\|I_0 \mathcal{T}(h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2$ and $\|h_{\lambda_{\mathcal{H}}}^\dagger - h_0^\dagger\|_{\mathcal{H}}^2$, we conclude that, with probability at least $\kappa(\delta_n)$, $\|\tilde{h} - h_0^\dagger\|_2^2 \leq c_2 \left(\frac{\delta_n^2}{\lambda_{\mathcal{H}}} + \frac{\lambda_{\mathcal{G}}}{\lambda_{\mathcal{H}}} + \lambda_{\mathcal{H}}^{\min\{\beta, 1\}} \right)$, where $c_2 = k_1^2 \cdot \max\{2c_1^* + c_2^* c^{-2}, 2B_1^2, 2c_0''\}$, and c_1^* and c_2^* are the constants of Theorem 4.3.

To derive the bound for $\|I_0 \mathcal{T}(\tilde{h} - h_0^\dagger)\|_2^2$, first note that from inequality (C.2) we have that the event

$$\frac{1}{8c^2} \|I_0 \mathcal{T}(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 \leq \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + \lambda_{\mathcal{G}} B_1^2 + c_0'' \delta_n^2$$

happens with probability at least $\kappa(\delta_n)$. Combining this inequality with the inequality

$\|I_0\mathcal{T}(\tilde{h} - h_0^\dagger)\|_2^2 \leq 2\|I_0\mathcal{T}(\tilde{h} - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + 2\|I_0\mathcal{T}(h_{\lambda_{\mathcal{H}}}^\dagger - h_0^\dagger)\|_2^2$ we conclude that

$$\|I_0\mathcal{T}(\tilde{h} - h_0^\dagger)\|_2^2 \leq 10\|I_0\mathcal{T}(h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2 + 16c^2B_1^2\lambda_{\mathcal{G}} + 16c^2c_0''\delta_n^2$$

holds with probability at least $\kappa(\delta_n)$. Invoking again Theorem 4.3 to bound $\|I_0\mathcal{T}(h_0^\dagger - h_{\lambda_{\mathcal{H}}}^\dagger)\|_2^2$, we conclude that the event $\|I_0\mathcal{T}(\tilde{h} - h_0^\dagger)\|_2^2 \leq c_3 \left(\delta_n^2 + \lambda_{\mathcal{G}} + \lambda_{\mathcal{H}}^{\min\{\beta+1, 2\}} \right)$ happens with probability at least $\kappa(\delta_n)$ where $c_3 = \max\{10c_2^*, 16c^2B_1^2, 16c^2c_0''\}$. This completes the proof. $\square$

Verifying Assumption 4.7 for the Map $\tilde{m}(o; h)$ from Example 4.2. We denote the support of random variables $(V_1, \dots, V_k)$ and their conditional versions $V_1|(V_2 = v_2, \dots, V_k = v_k)$, for $k \geq 2$, as $\text{supp}(V_1, \dots, V_k)$ and $\text{supp}(V_1|V_2 = v_2, \dots, V_k = v_k)$.

Observe that q , defined in Example 4.2, is strictly monotone and differentiable a.e. with respect to $P_{0,A}$, and admits an inverse that is differentiable a.e. with positive derivative.

Suppose that $\tilde{\rho}_{P_0}(a, l, w) = I_0(a, l) \cdot \frac{dq^{-1}(a)}{da} \cdot \frac{p_{0,A|L,W}\{q^{-1}(a)|l, w\}}{p_{0,A|L,W}(a|l, w)}$, where $I_0(a, l) := I\{a \in [c + \delta, d]\}$, is bounded by some $B_0 > 0$. Then, for any $h \in \mathcal{L}^2(A, L, W)$,

$$\begin{aligned} \|\tilde{m}(O; h)\|_2^2 &= \iint \left[\int h\{q(a), l, w\}^2 p_{0,A|L,W}(a|l, w) da \right] dP_0(l, w) \\ &\stackrel{(i)}{=} \iiint h(\tilde{a}, l, w)^2 \cdot I\{\tilde{a} \in [c + \delta, d]\} \cdot \frac{dq^{-1}(\tilde{a})}{d\tilde{a}} p_{0,A|L,W}(q^{-1}(\tilde{a})|l, w) d\tilde{a} \cdot dP_0(l, w) \\ &\stackrel{(ii)}{=} \iiint h(\tilde{a}, l, w)^2 \cdot I_0(\tilde{a}, l) \cdot \tilde{\rho}_{P_0}(\tilde{a}, l, w) p_{0,A|L,W}(\tilde{a}|l, w) d\tilde{a} \cdot dP_0(l, w) \\ &\stackrel{(iii)}{\leq} B_0 \iiint I_0(\tilde{a}, l) \cdot h(\tilde{a}, l, w)^2 p_{0,A|L,W}(\tilde{a}|l, w) d\tilde{a} \cdot dP_0(l, w) = B_0 \cdot \|I_0 h\|_2^2, \end{aligned}$$

where the equality in (i) follows from the change of variable $\tilde{a} = q(a)$, which is valid since q has an a.e. differentiable inverse with strictly positive derivative and maps $[c, d]$ onto $[c + \delta, d]$. Equality (ii) follows directly from the definition of $\tilde{\rho}_{P_0}$, and the inequality (iii) uses

the assumption that $|\tilde{\rho}_{P_0}(a, l, w)| \leq B_0$ for all $(a, l, w) \in \text{supp}(A, L, W)$. This proves the first inequality in part 1 of Assumption 4.7 with constant $B_2 = \sqrt{B_0}$.

To verify part 2 of Assumption 4.7, we impose the following *common support* assumption: $\text{supp}(A|L = l, U = u) = \text{supp}(A|L = l)$ for all $(l, u) \in \text{supp}(L, U)$, which is standard in proximal inference (see Section 3.3 of Chapter 3). This assumption, together with the fact that W is a negative control outcome—i.e., $(A, Z) \perp W|L, U$ —imply that $\text{supp}(A|L = l, W = w) = \text{supp}(A|L = l)$ for all $(l, w) \in \text{supp}(L, W)$. Moreover, $\text{supp}(A, L, W) = [c, d] \times \text{supp}(L, W)$ given that $\text{supp}(A|L = l) = [c, d]$ for all $l \in \mathcal{L}$.

Next, we define an equivalence relation on $\mathcal{H}$ as follows: $h \leftrightarrow h'$ if and only if $h(a, l, w) = h'(a, l, w)$ for all $(a, l, w) \in [c + \delta, d] \times \text{supp}(L, W)$. Let $\widetilde{\mathcal{H}} := \{[h] : h \in \mathcal{H}\}$ denote the set of equivalence classes in $\mathcal{H}$ under this relation. Define the map $d : \widetilde{\mathcal{H}} \times \widetilde{\mathcal{H}} \mapsto \mathbb{R}$ by

$$\tilde{d}([h_1], [h_2]) := \inf\{\|h'_1 - h'_2\|_{\mathcal{H}} \mid h'_1 \in [h_1], h'_2 \in [h_2]\}.$$

We now verify that d is a metric on $\widetilde{\mathcal{H}}$. By construction, d is symmetric. To see that the triangle inequality holds, take $[h_1], [h_2], [h_3] \in \widetilde{\mathcal{H}}$, and any representatives $h_1 \in [h_1]$, $h_2 \in [h_2]$, and $h_3 \in [h_3]$. Then, $\|h_1 - h_2\|_{\mathcal{H}} \leq \|h_1 - h_3\|_{\mathcal{H}} + \|h_3 - h_2\|_{\mathcal{H}}$. Taking the infimum over representatives on both sides yields $d([h_1], [h_2]) \leq d([h_1], [h_3]) + d([h_3], [h_2])$, so d satisfies the triangle inequality.

Finally, we show that $d([h_1], [h_2]) = 0$ if and only if $[h_1] = [h_2]$. The if direction is trivial since $d([h], [h]) = 0$ for all $[h] \in \widetilde{\mathcal{H}}$. To show the only if direction suppose that $d([h_1], [h_2]) = 0$ but $[h_1] \neq [h_2]$. Then there exist representatives $h_1 \in [h_1]$, $h_2 \in [h_2]$, and a point $(a_0, l_0, w_0) \in [c + \delta, d] \times \text{supp}(L, W)$ such that $h_1(a_0, l_0, w_0) \neq h_2(a_0, l_0, w_0)$. Let $\epsilon := |h_1(a_0, l_0, w_0) - h_2(a_0, l_0, w_0)| > 0$, and $k := [\sup_{(a, l, w) \in \text{supp}(A, L, W)} K_{\mathcal{H}}\{(a, l, w), (a, l, w)\}]^{1/2} < \infty$. By the definition of d and the fact that $d([h_1], [h_2]) = 0$, there exists $h'_1 \in [h_1]$ and

$h'_2 \in [h_2]$ such that $\|h'_1 - h'_2\|_{\mathcal{H}} < \epsilon/k$. Then we arrive at the following contradiction:

$$k\|h'_1 - h'_2\|_{\mathcal{H}} < |h_1(a_0, l_0, w_0) - h_2(a_0, l_0, w_0)| = |h'_1(a_0, l_0, w_0) - h'_2(a_0, l_0, w_0)| \leq k\|h'_1 - h'_2\|_{\mathcal{H}},$$

where the equality holds because $h_1 \leftrightarrow h'_1$ and $h_2 \leftrightarrow h'_2$, and the final inequality because for any $h \in \mathcal{H}$, $\|h\|_{\infty} \leq k \|h\|_{\mathcal{H}}$. Hence, $d([h_1], [h_2]) = 0$ implies $[h_1] = [h_2]$, so d is a valid metric on $\widetilde{\mathcal{H}}$.

We now show that $(\widetilde{\mathcal{H}}, d)$ is separable. Since $K_{\mathcal{H}}$ is continuous and $\text{supp}(A, L, W)$ is compact, $\mathcal{H}$ is separable under the metric induced by its norm (see Proposition 11.7 of [Paulsen and Raghupathi \(2016\)](#)). Let $\{h_n\} \subset \mathcal{H}$ be a countable dense subset, and consider the set of equivalence classes $\{[h_n]\}$. We claim this set is dense in $\widetilde{\mathcal{H}}$. Indeed, for any $[h] \in \widetilde{\mathcal{H}}$ and $\epsilon > 0$, there exists a representative $h' \in [h]$ and some $h_n \in \mathcal{H}$ such that $\|h' - h_n\|_{\mathcal{H}} < \epsilon$. Then $d([h], [h_n]) \leq \|h - h_n\|_{\mathcal{H}} < \epsilon$, so $\{[h_n]\}$ is dense in $\widetilde{\mathcal{H}}$, and thus $\widetilde{\mathcal{H}}$ is separable.

Next, consider the class $\widetilde{m} \circ \mathcal{H} := \{o \mapsto h(q(a), l, w)\}$, which is in a one-to-one correspondence with $\widetilde{\mathcal{H}}$. More precisely, the map $h \mapsto \widetilde{m}(\cdot, h)$ is well-defined on equivalence classes, since

$$\begin{aligned} & \widetilde{m}(o; h_1) = \widetilde{m}(o; h_2) \quad \forall o \in \mathcal{O} \\ \iff & h_1(q(a), l, w) = h_2(q(a), l, w) \quad \forall (a, l, w) \in [c, d] \times \text{supp}(L, W) \\ \stackrel{(i)}{\iff} & h_1(a, l, w) = h_2(a, l, w) \quad \forall (a, l, w) \in [c + \delta, d] \times \text{supp}(L, W) \\ \iff & h_1 \leftrightarrow h_2, \end{aligned}$$

where (i) follows because $q([c, d]) = [c + \delta, d]$. Thus, we can equip $\widetilde{m} \circ \mathcal{G}$ with the metric d defined as:

$$d(f_1, f_2) := \widetilde{d}([h_{f_1}], [h_{f_2}]) \quad \text{for any } f_1, f_2 \in \widetilde{\mathcal{H}},$$

where $[h_{f_1}]$ and $[h_{f_2}]$ are their corresponding equivalence classes in $\widetilde{\mathcal{H}}$. This yields $(\widetilde{m} \circ \mathcal{H}, d)$ a separable metric space. By construction, the metric d is controlled by the RKHS norm.

Finally, we show that the metric d dominates the sup-norm on $\widetilde{m} \circ \mathcal{H}$. For any $f_1, f_2 \in \widetilde{m} \circ \mathcal{H}$, let $[h_{f_1}]$ and $[h_{f_2}]$ be their corresponding equivalence classes, and take any representatives $h_{f_1} \in [h_{f_1}]$ and $h_{f_2} \in [h_{f_2}]$. Then, for any $o \in \mathcal{O}$, we have

$$|f_1(o) - f_2(o)| := |\widetilde{m}(o; h_{f_1}) - \widetilde{m}(o; h_{f_2})| = |h_{f_1}(q(a), l, w) - h_{f_2}(a, l, w)| \leq k \|h_{f_1} - h_{f_2}\|_{\mathcal{H}}.$$

Since the representatives were arbitrary, we conclude that $\sup_{o \in \mathcal{O}} |f_1(o) - f_2(o)| \leq k \cdot \widetilde{d}([h_1], [h_2]) = k \cdot d(f_1, f_2)$, implying that d dominates the metric induced by the sup-norm with constant $B_3 = k$. This completes the verification of part 2 of Assumption 4.7. $\square$

Theorem 4.5. The proof follows analogously to that of Theorem 4.4, but applies Lemma C.11 in place of Lemma C.10, and uses the updated constants defined in Lemma C.11. $\square$

Lemma 4.1. Let $P_0^q := P_0 \circ q^{-1}$ be the pushforward measure of P_0 under q . For any $g \in \mathcal{G}_{B_1}$, the composition $m(\cdot, g) = g \circ q(\cdot)$ satisfies

$$\|m(\cdot, g)\|_{\mathcal{L}^2(P_0)}^2 = \|g \circ q\|_{\mathcal{L}^2(P_0)}^2 = \int_{\mathcal{Z}} g\{q(z)\}^2 dP_0(z) = \int_{q(\mathcal{Z})} g(\tilde{z})^2 d(P_0 \circ q^{-1})(\tilde{z}) = \|g\|_{\mathcal{L}^2(P_0^q)}^2,$$

so the map $g \mapsto g \circ q = m(\cdot, g)$ is an isometry between $(\mathcal{G}_{B_1}, \mathcal{L}^2(P_0^q))$ and $(m \circ \mathcal{G}_{B_1}, \mathcal{L}^2(P_0))$.

Observe that the Rademacher complexity is preserved under reparameterization: if

$Z_1, \dots, Z_n \sim P_0$, $\tilde{Z}_1, \dots, \tilde{Z}_n \sim P_0^q$, and $\varepsilon_1, \dots, \varepsilon_n$ are i.i.d. Rademacher variables, then $\frac{1}{n} \sum_{i=1}^n \varepsilon_i g\{q(Z_i)\}^2 = \frac{1}{n} \sum_{i=1}^n \varepsilon_i g(\tilde{Z}_i)^2$, so for all $\delta > 0$, $\mathcal{R}_n(m \circ \mathcal{G}_{B_1}, P_0, \delta) = \mathcal{R}_n(\mathcal{G}_{B_1}, P_0^q, \delta)$. Since $o \mapsto m(o, g) = g(q(z))$ is linear in g and $\mathcal{G}_{B_1}$ is star-shaped, $m \circ \mathcal{G}_{B_1}$ is also star-shaped. Thus, $\mathcal{R}_n(\text{star}(m \circ \mathcal{G}_{B_1}), P_0, \delta) = \mathcal{R}_n(\mathcal{G}_{B_1}, P_0^q, \delta)$. Then the critical radius of $m \circ \mathcal{G}_{B_1}$ under P_0 equals the critical radius of $\mathcal{G}_{B_1}$ under P_0^q . The result follows because the critical radius

of an RKHS class is preserved up to multiplicative constants under changes of measure. $\square$

Ensuring Bounds in Theorems 4.4 and 4.5 are Achieved with High Probability. When $\delta_n < 5b/6$, we show that $\log(\log(b/\delta_n)) \leq c_1 n \delta_n^2 / 2$, where c_1 is the constant defined in Theorem 4.4, for a suitable choice of δ_n . Suppose $\delta_n^2 \geq c_0 \log(\log(n))/n$ with $c_0 = \max\{b^2/\log(\log(3)), 2/c_1\}$. Since $\delta_n/b < 1$, it holds that $\delta_n/b \geq (\delta_n/b)^2 \geq c_0 \log(\log(n))/(nb^2)$. If $n \geq 3$, the assumption $c_0 \geq b^2/\log(\log(3))$ implies $\delta_n/b \geq \log(\log(n))/\{n \log(\log(3))\} \geq 1/n$ which then implies $n \geq b/\delta_n$. Furthermore, since $c_0 \geq 2/c_1$, we have $\log(\log(b/\delta_n)) \leq \log(\log(n)) \leq n \delta_n^2 / c_0 \leq c_1 n \delta_n^2 / 2$, as desired. $\square$

Deriving the adjoint of $T_{\mathcal{G}} \circ \mathcal{T} : \mathcal{H} \rightarrow \mathcal{G}$. For $h \in \mathcal{H}$ and $g \in \mathcal{G}$, applying Lemma C.1(g) with the integral operators induced by $K_{\mathcal{H}}$ and $K_{\mathcal{G}}$ yields $\langle T_{\mathcal{G}} \circ \mathcal{T}h, g \rangle_{\mathcal{G}} = \langle \mathcal{T}h, g \rangle_{\mathcal{L}^2(P_Z)} = \langle h, \mathcal{T}^*g \rangle_{\mathcal{L}^2(P_W)} = \langle h, T_{\mathcal{H}} \circ \mathcal{T}^*g \rangle_{\mathcal{H}}$. Hence, the adjoint of $T_{\mathcal{G}} \circ \mathcal{T}$ is $T_{\mathcal{H}} \circ \mathcal{T}^*$. $\square$

C.2 Supporting Lemmas

The following Lemma establishes several important properties of the operator $T_{\mathcal{H}}^{1/2}$ that are invoked in the proofs of Lemmas C.2 and C.4, Theorem 4.3, and in the derivation of expression (4.15). Let $\mathcal{N}(T_{\mathcal{H}}) := \{h \in \mathcal{L}^2(P_{0,W}) : T_{\mathcal{H}}h = 0\}$ denote the null space of the operator $T_{\mathcal{H}}$, and let $\mathcal{N}(T_{\mathcal{H}})^{\perp}$ denote its orthogonal complement in $\mathcal{L}^2(P_{0,W})$.

Lemma C.1. *Suppose Assumption 4.2 holds. The following holds:*

- a) $\mathcal{N}(T_{\mathcal{H}}) = \mathcal{N}(T_{\mathcal{H}}^{1/2})$.
- b) For any $h \in \mathcal{L}^2(P_{0,W})$, it holds that $T_{\mathcal{H}}^{1/2}h \in \mathcal{H}$.
- c) The linear operator $T_{\mathcal{H}}^{1/2} : \mathcal{L}^2(P_{0,W}) \rightarrow \mathcal{H}$ is bounded and satisfies $\|\mathcal{T}_{\mathcal{H}}^{1/2}h\|_{\mathcal{H}} \leq \|h\|_2$.
- d) The restriction of $T_{\mathcal{H}}^{1/2}$ to $\mathcal{N}(T_{\mathcal{H}})^{\perp}$ is an isometry, i.e. $\|T_{\mathcal{H}}^{1/2}h\|_{\mathcal{H}} = \|h\|_2$ for all $h \in \mathcal{N}(T_{\mathcal{H}})^{\perp}$.

- e) The operator $T_{\mathcal{H}}^{1/2} : \mathcal{N}(T_{\mathcal{H}})^{\perp} \rightarrow \mathcal{H}$ is bijective.
- f) The operator $T_{\mathcal{H}}^{1/2} : \mathcal{L}^2(P_{0,W}) \rightarrow \mathcal{H}$, with its range endowed with the $\mathcal{L}^2(P_{0,W})$ -norm, is compact.
- g) For any $h' \in \mathcal{L}^2(P_W)$ and $h \in \mathcal{H}$, it holds that $\langle T_{\mathcal{H}} h', h \rangle_{\mathcal{H}} = \langle h', h \rangle_{\mathcal{L}^2(P_W)}$.

Proof. As discussed in Section 4.5, Assumption 4.2 implies that the integral operator $T_{\mathcal{H}} : \mathcal{L}^2(P_{0,W}) \mapsto \mathcal{L}^2(P_{0,W})$ is compact and positive. Therefore, by the spectral decomposition theorem (see, e.g., Theorem 4.3 of Hohage (2002)) there exists eigenfunctions $(\varphi_j)_{j=1}^N$ corresponding to all the non-zero eigenvalues $(\eta_j)_{j=1}^N$ of $T_{\mathcal{H}}$, where $N \in \mathbb{N} \cup \{+\infty\}$. Such eigenfunctions are orthonormal in $\mathcal{L}^2(P_{0,W})$ and satisfy that for any $h \in \mathcal{L}^2(P_{0,W})$, it holds that $T_{\mathcal{H}} h = \sum_{j=1}^N \eta_j \langle h, \varphi_j \rangle_2 \varphi_j$. Because the integral operator $T_{\mathcal{H}}$ is positive, its non-zero eigenvalues $(\eta_j)_{j=1}^N$ are positive. When $N = \infty$, $\eta_j \rightarrow 0$ as $j \rightarrow \infty$ and the sequence $(\eta_j)_{j=1}^N$ is arranged in decreasing order. The square root of $T_{\mathcal{H}}$, denoted by $T_{\mathcal{H}}^{1/2}$, is defined by $T_{\mathcal{H}}^{1/2} h := \sum_{j=1}^N \sqrt{\eta_j} \langle h, \varphi_j \rangle_2 \varphi_j$.

Assertion (a) follows from the fact that for any $h \in \mathcal{L}^2(P_{0,W})$: $T_{\mathcal{H}} h = \sum_{j=1}^N \eta_j \langle h, \varphi_j \rangle_2 \varphi_j = 0$ iff $\langle h, \varphi_j \rangle_2 = 0 \forall j = 1, \dots, N$ iff $T_{\mathcal{H}}^{1/2} h = \sum_{j=1}^N \sqrt{\eta_j} \langle h, \varphi_j \rangle_2 \varphi_j = 0$.

Another result (see Theorem 11.18 of Paulsen and Raghupathi (2016)), establishes that under Assumption 4.2, the sequence $(\sqrt{\eta_j} \varphi_j)_{j=1}^N$ forms an orthogonal basis of the RKHS $\mathcal{H}$. Furthermore, $\mathcal{H}$ admits the representation

$$\mathcal{H} = \left\{ h = \sum_{j=1}^N \gamma_j \varphi_j : \text{for some } (\gamma_j)_{j=1}^{\infty} \text{ with } \sum_{j=1}^N \frac{\gamma_j^2}{\eta_j} < \infty \right\}, \quad (\text{C.3})$$

where if $N = \infty$, the limit of the finite sums is with respect to the $\mathcal{L}^2(P_{0,W})$ -metric. Additionally, its inner product decomposes as (see Theorem 11.3 of Paulsen and Raghupathi (2016)) $\langle h, h' \rangle_{\mathcal{H}} = \sum_{j=1}^N \langle h, \varphi_j \rangle_2 \langle h', \varphi_j \rangle_2 / \eta_j$.

Assertion (b) follows from the representation of $\mathcal{H}$ in (C.3) and the definition of $T_{\mathcal{H}}^{1/2} h$.

Specifically, for any $h \in \mathcal{L}^2(P_{0,W})$, we have $T_{\mathcal{H}}^{1/2}h = \sum_{j=1}^N \sqrt{\eta_j} \langle h, \varphi_j \rangle_2 \varphi_j$, so to show that $T_{\mathcal{H}}^{1/2}h$ is in $\mathcal{H}$ we need to show that $\sum_{j=1}^N (\sqrt{\eta_j} \langle h, \varphi_j \rangle_2)^2 / \eta_j < \infty$. This holds because

$$\sum_{j=1}^N \frac{(\sqrt{\eta_j} \langle h, \varphi_j \rangle_2)^2}{\eta_j} = \sum_{j=1}^N \langle h, \varphi_j \rangle_2^2 \leq \|h\|_2^2 < \infty. \quad (\text{C.4})$$

Assertion (c) follows from (C.4) since $\|T_{\mathcal{H}}^{1/2}h\|_{\mathcal{H}}^2 = \sum_{j=1}^N (\sqrt{\eta_j} \langle h, \varphi_j \rangle_2)^2 / \eta_j$, thus proving that $T_{\mathcal{H}}^{1/2}$ is bounded with operator norm at most one. The linearity of $T_{\mathcal{H}}^{1/2}$ is immediate from the linearity of the inner products $\langle h, \varphi_j \rangle_2$, completing the proof of (c).

To show Assertions (d) and (e), we use that $\mathcal{N}(T_{\mathcal{H}})^{\perp} = \text{span}(\varphi_1, \dots, \varphi_N)$. Assertion (d) follows because, for $h \in \text{span}(\varphi_1, \dots, \varphi_N)$, the first inequality in (C.4) becomes an equality.

To prove Assertion (e), we first show that $T_{\mathcal{H}}^{1/2} : \mathcal{N}(T_{\mathcal{H}})^{\perp} \rightarrow \mathcal{H}$ is injective. Let $h_1, h_2 \in \mathcal{N}(T_{\mathcal{H}})^{\perp}$ be such that $T_{\mathcal{H}}^{1/2}h_1 = T_{\mathcal{H}}^{1/2}h_2$. Then, $h_1 - h_2 \in \mathcal{N}(T_{\mathcal{H}}^{1/2})$. By part (a), $\mathcal{N}(T_{\mathcal{H}}^{1/2}) = \mathcal{N}(T_{\mathcal{H}})$. Then, $h_1 - h_2 \in \mathcal{N}(T_{\mathcal{H}}) \cap \mathcal{N}(T_{\mathcal{H}})^{\perp} = \{0\}$. This shows injectivity. Surjectivity follows from the representation of $\mathcal{H}$ in (C.3). Specifically, any $h \in \mathcal{H}$ can be written as $h = \sum_{j=1}^N \gamma_j \varphi_j$ with $\sum_{j=1}^N \frac{\gamma_j^2}{\eta_j} < \infty$. Then $T_{\mathcal{H}}^{1/2}h' = h$ where $h' := \sum_{j=1}^N \gamma_j \varphi_j / \sqrt{\eta_j} \in \mathcal{N}(T_{\mathcal{H}})^{\perp}$.

Assertion (f) follows from the definition of $\mathcal{T}_{\mathcal{H}}^{1/2} : \mathcal{L}^2(P_{0,W}) \mapsto \mathcal{L}^2(P_{0,W})$, as it is the norm-limit of finite-rank operators (see Proposition 5.6.5 of Conway (2012)).

To prove assertion (g), observe that for any $h' \in \mathcal{L}^2(P_W)$ and $h \in \mathcal{H}$, it holds that $\langle T_{\mathcal{H}}h', h \rangle_{\mathcal{H}} = \sum_{j=1}^N \langle T_{\mathcal{H}}h', \varphi_j \rangle_2 \langle h, \varphi_j \rangle_2 / \eta_j = \sum_{j=1}^N \langle h', \varphi_j \rangle_2 \langle h, \varphi_j \rangle_2 = \langle h', \sum_{j=1}^N \langle h, \varphi_j \rangle_2 \varphi_j \rangle_2$. For every $h \in \mathcal{H}$, $h = \sum_{j=1}^N \langle h, \varphi_j \rangle_2 \varphi_j$. Hence, the last expression equals $\langle h', h \rangle_2$, as claimed. $\square$

Lemma C.2. *Suppose Assumption 4.2 holds. Then the operators $I_0^{op} \circ \mathcal{T} \mid_{\mathcal{H}} : \mathcal{H} \rightarrow \mathcal{L}^2(P_Z)$ and $I_0^{op} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2} : \mathcal{L}^2(P_W) \rightarrow \mathcal{L}^2(P_Z)$ are bounded and linear, and their adjoints $(I_0^{op} \circ \mathcal{T} \mid_{\mathcal{H}})^* : \mathcal{L}^2(P_Z) \rightarrow \mathcal{H}$ and $(I_0^{op} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2})^* : \mathcal{L}^2(P_Z) \rightarrow \mathcal{L}^2(P_W)$ are given by $T_{\mathcal{H}} \circ \mathcal{T}^* \circ I_0^{op}$ and $T_{\mathcal{H}}^{1/2} \circ \mathcal{T}^* \circ I_0^{op}$.*

Proof. Linearity is immediate from that of I_0^{op} , $\mathcal{T}$, and $T_{\mathcal{H}}^{1/2}$. Boundedness of $I_0^{op} \circ \mathcal{T} \mid_{\mathcal{H}}$ follows because for $h \in \mathcal{H}$, it holds that $\|I_0^{op} \mathcal{T} h\|_2^2 \leq \|r\|_{\infty}^2 \|h\|_2^2 \leq \|r\|_{\infty}^2 \|h\|_{\infty}^2 \leq k_1^2 \|r\|_{\infty}^2 \|h\|_{\mathcal{H}}^2$, where $k_1^2 = \sup_{w \in \mathcal{W}} K_{\mathcal{H}}(w, w) < \infty$. Boundedness of $I_0^{op} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2}$ follows because $T_{\mathcal{H}}^{1/2}$ (by Lemma C.1(f)), $\mathcal{T}$, and I_0^{op} are all bounded, and so is their composition.

To determine the adjoint of $I_0^{op} \circ \mathcal{T} \mid_{\mathcal{H}}$, let $h \in \mathcal{H}$ and $g \in \mathcal{L}^2(P_Z)$. Then, applying Lemma C.1(g) yields $\langle I_0^{op} \circ \mathcal{T} h, g \rangle_{\mathcal{L}^2(P_Z)} = \langle \mathcal{T} h, I_0^{op} g \rangle_{\mathcal{L}^2(P_Z)} = \langle h, \mathcal{T}^* \circ I_0^{op} g \rangle_{\mathcal{L}^2(P_W)} = \langle h, T_{\mathcal{H}} \circ \mathcal{T}^* \circ I_0^{op} g \rangle_{\mathcal{H}}$. Hence, $(I_0^{op} \circ \mathcal{T} \mid_{\mathcal{H}})^* = T_{\mathcal{H}} \circ \mathcal{T}^* \circ I_0^{op}$. Now let $h \in \mathcal{L}^2(P_W)$ and $g \in \mathcal{L}^2(P_Z)$. Since $T_{\mathcal{H}}^{1/2}$ is self-adjoint, we have $\langle I_0^{op} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2} h, g \rangle_{\mathcal{L}^2(P_Z)} = \langle T_{\mathcal{H}}^{1/2} h, \mathcal{T}^* \circ I_0^{op} g \rangle_{\mathcal{L}^2(P_W)} = \langle h, T_{\mathcal{H}}^{1/2} \circ \mathcal{T}^* \circ I_0^{op} g \rangle_{\mathcal{L}^2(P_W)}$. Thus $(I_0^{op} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2})^* = T_{\mathcal{H}}^{1/2} \circ \mathcal{T}^* \circ I_0^{op}$. This completes the proof. $\square$

Lemma C.3. *The second equality in display (4.15) holds.*

Proof. To show the second equality in (4.15) it suffices to show that $\|h_{\lambda_{\mathcal{H}}, KRAS}^{\dagger\dagger}\|_{\mathcal{H}} \leq B$ where $h_{\lambda_{\mathcal{H}}, KRAS}^{\dagger\dagger} := \arg \min_{h \in \mathcal{H}} \left\{ \frac{1}{4} \|\rho - I_0^{op} \circ \mathcal{T} \mid_{\mathcal{H}} h\|_2^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2 \right\}$. This holds because $\|h_{0, KRAS}^{\dagger}\|_{\mathcal{H}} < B$ by definition of B and $\|h_{\lambda_{\mathcal{H}}, KRAS}^{\dagger\dagger}\|_{\mathcal{H}} \leq \|h_{0, KRAS}^{\dagger}\|_{\mathcal{H}}$ because otherwise $\frac{1}{4c^2} \|I_0 \mathcal{T}(h_{0, KRAS}^{\dagger} - h_{\lambda_{\mathcal{H}}, KRAS}^{\dagger\dagger})\|_2^2 + \lambda_{\mathcal{H}} \|h_{0, KRAS}^{\dagger}\|_{\mathcal{H}}^2 < \frac{1}{4c^2} \|I_0 \mathcal{T}(h_{0, KRAS}^{\dagger} - h_{\lambda_{\mathcal{H}}, KRAS}^{\dagger\dagger})\|_2^2 + \lambda_{\mathcal{H}} \|h_{\lambda_{\mathcal{H}}, KRAS}^{\dagger\dagger}\|_{\mathcal{H}}^2$ violating the definition of $h_{\lambda_{\mathcal{H}}, KRAS}^{\dagger\dagger}$. $\square$

Lemma C.4. *Under Assumptions 4.1 and 4.2, $h_{0, KRAS}^{\dagger} = T_{\mathcal{H}}^{1/2} h'_{0, KRAS}$ and $h'_{0, KRAS} \in \mathcal{N}(T_{\mathcal{H}})^{\perp}$.*

Proof. To simplify notation, we drop the subscripts $KRAS$ from $h_{0, KRAS}^{\dagger}$ and $h'_{0, KRAS}$. Under Assumptions 4.1–4.2, $h_0^{\dagger}$ in (4.13) is well-defined, and the $\widetilde{\mathcal{T}}$ is bounded (see Section 4.5).

By Lemma C.1 (e), there exists a unique $h'_0 \in \mathcal{N}(T_{\mathcal{H}})^{\perp}$ such that $h_0^{\dagger} = T_{\mathcal{H}}^{1/2} h'_0$. Since $\rho = \mathcal{T}(I_0 h_0^{\dagger}) = I_0^{op} \circ \mathcal{T} h_0^{\dagger} = I_0^{op} \circ \mathcal{T} \circ Id_{\mathcal{H}} h_0^{\dagger} = I_0^{op} \circ \mathcal{T} \circ Id_{\mathcal{H}} \circ T_{\mathcal{H}}^{1/2} h'_0 = \widetilde{\mathcal{T}} h'_0$ it follows that h'_0 solves equation (4.17). It remains to show that it is the solution of minimal $\mathcal{L}^2(P_W)$ –norm. Let $h_1 \in \mathcal{L}^2(P_W)$ be any solution to (4.17). Decompose h_1 as $h'_1 + h_1^{\perp}$, where $h'_1 \in \mathcal{N}(T_{\mathcal{H}})$ and

$h_1^\perp \in \mathcal{N}(T_{\mathcal{H}})^\perp$. Since $\widetilde{\mathcal{T}}h_1^\perp = \widetilde{\mathcal{T}}h_1$, it follows that $h_1^\perp$ also solves equation (4.17). Moreover, $\rho = \widetilde{\mathcal{T}}h_1^\perp = I_0^{op} \circ \mathcal{T} \circ Id_{\mathcal{H}} \circ T_{\mathcal{H}}^{1/2}h_1^\perp = I_0^{op} \circ \mathcal{T} \circ T_{\mathcal{H}}^{1/2}h_1^\perp = \mathcal{T}(I_0T_{\mathcal{H}}^{1/2}h_1^\perp)$, so $T_{\mathcal{H}}^{1/2}h_1^\perp \in \mathcal{H}_0$. By Lemma C.1 (d), $T_{\mathcal{H}}^{1/2}$ restricted to $\mathcal{N}(T_{\mathcal{H}})^\perp$ is an isometry, so $\|h_1^\perp\|_2 = \|T_{\mathcal{H}}^{1/2}h_1^\perp\|_{\mathcal{H}} \geq \|h_0^\dagger\|_{\mathcal{H}}^2 = \|h'_0\|_2^2$ with equality iff $T_{\mathcal{H}}^{1/2}h_1^\perp = h_0^\dagger = T_{\mathcal{H}}^{1/2}h'_0$, which by Lemma C.1 (e) implies $h_1^\perp = h'_0$. Therefore, $\|h_1\|_2^2 = \|h'_1\|_2^2 + \|h_1^\perp\|_2^2 \geq \|h'_0\|_2^2$, with equality, i.e., $\|h_1\|_2 = \|h'_0\|_2$, if and only if $h'_1 = 0$ and $h_1^\perp = h'_0$. This shows that h'_0 is the unique solution to equation (4.17) of minimal $\mathcal{L}^2(P_W)$ -norm. $\square$

Next we will present several results that are used in the proofs of Theorems 4.4 and 4.5.

In the following Lemmas C.5–C.9, $O_1, \dots, O_n$ are n independent identically distributed copies of a random variable O on the sample space $\mathcal{O}$. Furthermore, $\mathcal{F}$ is an index set such that for each $f \in \mathcal{F}$, $t_f : \mathcal{O} \mapsto \mathbb{R}$ is a measurable, real-valued function satisfying $\mathbb{E}[t_f(O)^2] < \infty$.

Lemma C.5 (Talagrand inequality, Theorem 3.27 of Wainwright (2019)). *Suppose $\mathcal{F}$ is countable and $\{t_f : f \in \mathcal{F}\}$ is uniformly bounded by b . Define $Z := \sup_{f \in \mathcal{F}} \left\{ \frac{1}{n} \sum_{i=1}^n t_f(O_i) \right\}$. Then, for all $u > 0$, $P\{Z \geq \mathbb{E}(Z) + u\} \leq 2 \exp\left(-\frac{nu^2}{16eb\mathbb{E}(Z) + 8e \sup_{f \in \mathcal{F}} \mathbb{E}\{t_f(O)^2\} + 4bu}\right)$ where $e := \exp(1)$.*

The proof of Lemma C.5 is given in Wainwright (2019) under the assumption that $\mathcal{F}$ is countable. However, the proofs of Theorems 4.4 and 4.5 require applying the lemma to uncountable sets $\mathcal{F}$. This extension is valid when $\{t_f : f \in \mathcal{F}\}$ are Caratheodory sets, i.e. when $\mathcal{F}$ can be equipped with a separable metric under which the evaluation map $f \mapsto t_f(o)$ is continuous for each $o \in \mathcal{O}$, a condition satisfied by the function classes we consider. For completeness, we now state and prove some results that imply that the Lemma C.5 remains valid under this more general setting.

To facilitate the extension of Lemma C.5 to certain uncountable uniformly bounded

function classes $\{t_f : f \in \mathcal{F}\}$ used in the proof of Theorem 4.4, in the next Lemma we adapt a result from [Steinwart and Christmann \(2008\)](#).

Lemma C.6. *Suppose that:*

1. $\mathcal{F}$ is a separable metric space with metric d , and let $\mathcal{S} \subset \mathcal{F}$ be a countable dense set,
2. For each $f \in \mathcal{F}$, $t_f : \mathcal{O} \mapsto \mathbb{R}$ is a measurable function, where $\mathcal{O}$ is the sample space of a random vector O ,
3. For each $o \in \mathcal{O}$, the map $f \mapsto t_f(o)$ is continuous with respect to d at all $f \in \mathcal{F}$.

Then for all $o_1, \dots, o_n \in \mathcal{O}$, $\sup_{f \in \mathcal{F}} \left\{ \frac{1}{n} \sum_{i=1}^n t_f(o_i) \right\} = \sup_{f \in \mathcal{S}} \left\{ \frac{1}{n} \sum_{i=1}^n t_f(o_i) \right\}$ and, consequently, for $O_1, \dots, O_n$ i.i.d. copies of O , $\sup_{f \in \mathcal{F}} \left\{ \frac{1}{n} \sum_{i=1}^n t_f(O_i) \right\}$ is a well-defined random variable.

Proof. Fix $\underset{\sim}{o} := (o_1, \dots, o_n) \in \mathcal{O}^n$ and define $g(\underset{\sim}{o}; f) := \frac{1}{n} \sum_{i=1}^n t_f(o_i)$, $A_1(\underset{\sim}{o}) := \sup_{f \in \mathcal{F}} g(\underset{\sim}{o}; f)$, and $A_2(\underset{\sim}{o}) := \sup_{f \in \mathcal{S}} g(\underset{\sim}{o}; f)$. We will show that $A_1(\underset{\sim}{o}) = A_2(\underset{\sim}{o})$.

Clearly $A_2(\underset{\sim}{o}) \leq A_1(\underset{\sim}{o})$ because $\mathcal{S} \subset \mathcal{F}$. Suppose that $A_2(\underset{\sim}{o}) < A_1(\underset{\sim}{o})$ and set $\varepsilon := A_1(\underset{\sim}{o}) - A_2(\underset{\sim}{o}) > 0$. By the definition of the supremum, there exists $f_\varepsilon \in \mathcal{F}$ such that $A_1(\underset{\sim}{o}) - \frac{\varepsilon}{2} < g(\underset{\sim}{o}; f_\varepsilon)$. Now, the map $f \mapsto g(\underset{\sim}{o}; f)$ is continuous because the map $f \mapsto t_f(o_i)$ is continuous for each i . Then, there exists $\delta > 0$ such that $d(f, f_\varepsilon) < \delta$ implies $|g(\underset{\sim}{o}; f) - g(\underset{\sim}{o}; f_\varepsilon)| < \frac{\varepsilon}{2}$. Since $\mathcal{S}$ is dense in $\mathcal{F}$, there exists $f_s \in \mathcal{S}$ such that $d(f_s, f_\varepsilon) < \delta$. Then, $g(\underset{\sim}{o}; f_s) > g(\underset{\sim}{o}; f_\varepsilon) - \frac{\varepsilon}{2} > A_1(\underset{\sim}{o}) - \varepsilon = A_2(\underset{\sim}{o})$, which contradicts the definition of $A_2(\underset{\sim}{o})$. We thus conclude that $A_1(\underset{\sim}{o}) = A_2(\underset{\sim}{o})$.

Next, since $g(\cdot; f) : \mathcal{O}^n \mapsto \mathbb{R}$ is measurable for all $f \in \mathcal{S}$, and the map $f \mapsto g(\underset{\sim}{o}; f)$ is continuous for all $\underset{\sim}{o} \in \mathcal{O}^n$, Lemma A.3.17 of [Steinwart and Christmann \(2008\)](#) ensures that the mapping $(\underset{\sim}{o}, f) \mapsto g(\underset{\sim}{o}; f)$ is measurable on $\mathcal{O}^n \times \mathcal{S}$. Then, the map $\underset{\sim}{o} \mapsto \sup_{f \in \mathcal{S}} g(\underset{\sim}{o}; f)$ is also measurable because $\mathcal{S}$ is countable. Hence, $\sup_{f \in \mathcal{S}} g(O_1, \dots, O_n; f) = \sup_{f \in \mathcal{F}} g(O_1, \dots, O_n; f)$ is a well-defined random variable. This completes the proof. $\square$

Corollary C.1. *Let $\mathcal{O}_0 \subset \mathcal{O}$ be a measurable set. Suppose that:*

1. $\mathcal{F}$ is a separable metric space with metric denoted by d such that $\mathcal{F} \subset \mathcal{L}^2(P_O)$, and let $\mathcal{S}$ denote a countable dense subset of it,
2. For each $f \in \mathcal{L}^2(P_O)$, $t_f : \mathcal{O} \mapsto \mathbb{R}$ is a measurable function, where $\mathcal{O}$ is the sample space of a random vector O ,
3. For each $o \in \mathcal{O}$, the map $f \mapsto t_{I_{\mathcal{O}_0} \cdot f}(o)$ is continuous with respect to d at all $f \in \mathcal{F}$.
4. The map $f \mapsto t_{I_{\mathcal{O}_0} \cdot f}$ is continuous from $(\mathcal{F}, d)$ into $\mathcal{L}^1(P_O)$.

For given fixed $\beta \in \mathbb{R}$ and any $o \in \mathcal{O}$, define the map $f \in \mathcal{F} \mapsto h_f(o) := t_{I_{\mathcal{O}_0} \cdot f}(o) - \beta \cdot \mathbb{E}\{t_{I_{\mathcal{O}_0} \cdot f}(O)\}$. Given $\mathcal{F}' \subset \mathcal{F}$, define $\mathcal{S}' := \mathcal{F}' \cap \mathcal{S}$. Then,

- i) $\sup_{f \in \mathcal{F}'} \left| \frac{1}{n} \sum_{i=1}^n h_f(O_i) \right| = \sup_{f \in \mathcal{S}'} \left| \frac{1}{n} \sum_{i=1}^n h_f(O_i) \right|$. In particular, the left-hand side is a well-defined random variable.
- ii) When $\beta = 0$, the conclusion in i) holds even if assumption 4 fails.

Proof. Fix $o \in \mathcal{O}$. First, we show that the map $f \mapsto h_f(o)$ is continuous with respect to the metric d . Take $f_1, f_2 \in \mathcal{F}$. By triangle inequality, we have $|h_{f_1}(o) - h_{f_2}(o)| \leq |t_{I_{\mathcal{O}_0} \cdot f_1}(o) - t_{I_{\mathcal{O}_0} \cdot f_2}(o)| + |\beta| \cdot \mathbb{E}_P\{|t_{I_{\mathcal{O}_0} \cdot f_1}(O) - t_{I_{\mathcal{O}_0} \cdot f_2}(O)|\}$.

Let $\varepsilon > 0$. By assumption 3, there exists $\delta > 0$ such that if $d(f_1, f_2) < \delta$, then $|t_{I_{\mathcal{O}_0} \cdot f_1}(o) - t_{I_{\mathcal{O}_0} \cdot f_2}(o)| < \varepsilon/2$. By assumption 4, there exists $\delta' > 0$ such that if $d(f_1, f_2) < \delta'$, then $|\beta| \cdot \mathbb{E}_P\{|t_{I_{\mathcal{O}_0} \cdot f_1}(O) - t_{I_{\mathcal{O}_0} \cdot f_2}(O)|\} < \varepsilon/2$. Then, if $d(f_1, f_2) < \min\{\delta, \delta'\}$, we have $|h_{f_1}(o) - h_{f_2}(o)| < \varepsilon$.

By Assumption 1 and definition of $\mathcal{F}'$, $\mathcal{F}'$ is separable with respect to d with dense subset $\mathcal{S}'$. Then, it follows from Lemma C.6 that

$$\sup_{f \in \mathcal{F}'} \left\{ \frac{1}{n} \sum_{i=1}^n h_f(O_i) \right\} = \sup_{f \in \mathcal{S}'} \left\{ \frac{1}{n} \sum_{i=1}^n h_f(O_i) \right\}, \quad \sup_{f \in \mathcal{F}'} \left\{ -\frac{1}{n} \sum_{i=1}^n h_f(O_i) \right\} = \sup_{f \in \mathcal{S}'} \left\{ -\frac{1}{n} \sum_{i=1}^n h_f(O_i) \right\},$$

and the left hand sides of both equalities are well-defined random variables. Therefore

$$\begin{aligned} \sup_{f \in \mathcal{F}'} \left| \frac{1}{n} \sum_{i=1}^n h_f(O_i) \right| &= \max \left\{ \sup_{f \in \mathcal{F}'} \left[\frac{1}{n} \sum_{i=1}^n h_f(O_i) \right], \sup_{f \in \mathcal{F}'} \left[-\frac{1}{n} \sum_{i=1}^n h_f(O_i) \right] \right\} \\ &= \max \left\{ \sup_{f \in \mathcal{S}'} \left[\frac{1}{n} \sum_{i=1}^n h_f(O_i) \right], \sup_{f \in \mathcal{S}'} \left[-\frac{1}{n} \sum_{i=1}^n h_f(O_i) \right] \right\} = \sup_{f \in \mathcal{S}'} \left| \frac{1}{n} \sum_{i=1}^n h_f(O_i) \right| \end{aligned}$$

and the leftmost hand side is a well-defined random variable. This concludes the proof. $\square$

Lemmas C.7–C.9 below are minor variations of results in [Wainwright \(2019\)](#) and [Foster and Syrgkanis \(2023\)](#). The following Lemma provides a weaker bound than Lemma 14.21 of [Wainwright \(2019\)](#).

Lemma C.7. *Suppose that:*

1. *The assumptions of Corollary C.1 hold with $\mathcal{F}$ satisfying: (1.1) $\mathcal{F}$ is uniformly bounded by b , (1.2) for each $o \in \mathcal{O}$, the evaluation functional $f \mapsto f(o)$ is continuous relative to d , and (1.3) $f = 0$ belongs to $\mathcal{F}$.*
2. *t_f satisfies $t_f = 0$ when $f = 0$.*
3. *There exists a constant L such that for any $f_1, f_2 \in \mathcal{L}^2(P_O)$, $|t_{f_1}(o) - t_{f_2}(o)| \leq L \cdot |f_1(o) - f_2(o)|$ for all $o \in \mathcal{O}$ ($o \in \mathcal{O}_0$).*

Define h_f as in Corollary C.1 with $\beta = 1$. For any $r > 0$ define $\mathcal{F}_r := \{f \in \mathcal{F} : \|I_{\mathcal{O}_0} \cdot f\|_2 \leq r\}$ and $Z_n(r) := \sup_{f \in \mathcal{F}_r} \left| \frac{1}{n} \sum_{i=1}^n h_f(O_i) \right|$. Then, $Z_n(r)$ is a well-defined random variable and for all $u > 0$, $P\{Z_n(r) \geq 4L\mathcal{R}_n(I_{\mathcal{O}_0} \cdot \mathcal{F}, r) + u\} \leq 2 \exp\left(-\frac{nu^2}{8eL^2r^2 + 128ebL^2\mathcal{R}_n(I_{\mathcal{O}_0} \cdot \mathcal{F}, r) + 8bLu}\right)$, where $e = \exp(1)$ and $\mathcal{R}_n(I_{\mathcal{O}_0} \cdot \mathcal{F}, r)$ is the localized Rademacher complexity of radius r of the function class $I_{\mathcal{O}_0} \cdot \mathcal{F} := \{I_{\mathcal{O}_0} \cdot f \mid f \in \mathcal{F}\}$. Moreover, if δ_n is any solution to the inequality $\mathcal{R}_n(\text{star}(I_{\mathcal{O}_0} \cdot \mathcal{F}), \delta) \leq \frac{\delta^2}{b}$, then for each $r \geq \delta_n$, $P\{Z_n(r) \geq \frac{4Lr\delta_n}{b} + u\} \leq 2 \exp\left(-\frac{nu^2}{8eL^2r^2 + 128eL^2r\delta_n + 8bLu}\right)$.

Proof. By assumption 1, the conclusion of Corollary C.1 holds; i.e., $Z_n(r)$ is well-defined.

Furthermore,

$$Z_n(r) = \sup_{f \in \mathcal{S}_r} \left| \frac{1}{n} \sum_{i=1}^n h_f(O_i) \right|$$

where $\mathcal{S}_r := \mathcal{S} \cap \mathcal{F}_r$. Now define the function class $\mathcal{C} := \{\alpha h_f : \alpha \in \{-1, 1\}, f \in \mathcal{S}_r\}$ which is countable by construction. With $\mathcal{C}$ so defined we have that $Z_n(r) = \sup_{h \in \mathcal{C}} \frac{1}{n} \sum_{i=1}^n h(O_i)$.

For any $h \in \mathcal{C}$, write $h = \alpha h_f$ for some $\alpha \in \{-1, 1\}$ and $f \in \mathcal{S}_r$. Letting $f' = 0$, we have

$$\begin{aligned} |h(o)| &\leq |t_{I_{\mathcal{O}_0} \cdot f}(o)| + \mathbb{E}\{|t_{I_{\mathcal{O}_0} \cdot f}(O)|\} \\ &= |t_{I_{\mathcal{O}_0} \cdot f}(o) - t_{I_{\mathcal{O}_0} \cdot f'}(o)| + \mathbb{E}\{|t_{I_{\mathcal{O}_0} \cdot f}(O) - t_{I_{\mathcal{O}_0} \cdot f'}(O)|\} \\ &\leq L \cdot I_{\mathcal{O}_0}(o) |f(o) - f'(o)| + L \cdot \mathbb{E}\{|I_{\mathcal{O}_0}(O)|f(O) - f'(O)|\} \leq 2Lb, \end{aligned}$$

where the first inequality is by the triangle inequality, the first and second equalities are by assumptions 1.3 and 2, the second inequality is by assumption 3, and the third inequality is by assumption 1.1. We thus conclude that $\mathcal{C}$ is uniformly bounded by $2Lb$. Since $\mathcal{C}$ is countable, applying Talagrand inequality (Lemma C.5) we obtain that, for any $u > 0$,

$$\mathbb{P}[Z_n(r) \geq \mathbb{E}\{Z_n(r)\} + u] \leq 2 \exp\left(-\frac{nu^2}{8e \sup_{h \in \mathcal{C}} \mathbb{E}[h(O)^2] + 32ebL\mathbb{E}\{Z_n(r)\} + 8bLu}\right).$$

Again, take $h \in \mathcal{C}$ and write $h = \alpha h_f$ for some $\alpha \in \{-1, 1\}$ and $f \in \mathcal{S}_r$. Thus $\mathbb{E}[h(O)^2] = \text{var}[t_{I_{\mathcal{O}_0} \cdot f}(O)] \leq \mathbb{E}[\{t_{I_{\mathcal{O}_0} \cdot f}(O)\}^2] \leq L^2 \mathbb{E}[\{I_{\mathcal{O}_0}(O) \cdot f(O)\}^2] \leq L^2 r^2$, where the second inequality follows by invoking assumption 3 applied to $f_1 = I_{\mathcal{O}_0} \cdot f$ and $f_2 = 0$ and noting that by assumption 2, $t_{f_2} = 0$, and the third inequality is because $f \in \mathcal{F}_r$.

We now bound $\mathbb{E}\{Z_n(r)\} = \mathbb{E}\{\sup_{f \in \mathcal{S}_r} |\frac{1}{n} \sum_{i=1}^n [t_{I_{\mathcal{O}_0} \cdot f}(O_i) - \mathbb{E}\{t_{I_{\mathcal{O}_0} \cdot f}(O)\}]|\}$ using a symmetrization argument. Letting $\{O'_i\}_{i=1}^n$ be independent copies of O and independent of $\{O_i\}_{i=1}^n$, we can write $\mathbb{E}\{Z_n(r)\} = \mathbb{E}\{\sup_{f \in \mathcal{S}_r} |\mathbb{E}[\frac{1}{n} \sum_{i=1}^n \{t_{I_{\mathcal{O}_0} \cdot f}(O_i) - t_{I_{\mathcal{O}_0} \cdot f}(O'_i)\}]|\}$. By the inequality $\sup_{d \in \mathcal{D}} \mathbb{E}\{|d(O)|\} \leq \mathbb{E}\{\sup_{d \in \mathcal{D}} |d(O)|\}$, which is valid for any countable class $\mathcal{D}$

of real-valued measurable functions, we conclude that

$$\mathbb{E}\{Z_n(r)\} \leq \mathbb{E}\left[\sup_{f \in \mathcal{S}_r} \left| \frac{1}{n} \sum_{i=1}^n \left\{ t_{I_{\mathcal{O}_0} \cdot f}(O_i) - t_{I_{\mathcal{O}_0} \cdot f}(O'_i) \right\} \right| \right].$$

Now, let $\varepsilon_1, \dots, \varepsilon_n$ be n i.i.d. Rademacher random variables, independent of O_i and O'_i for all $i = 1, \dots, n$. Given this independence assumption, for any function $f \in \mathcal{F}$, the $n \times 1$ random vector with i^{th} component $\varepsilon_i \{t_{I_{\mathcal{O}_0} \cdot f}(O_i) - t_{I_{\mathcal{O}_0} \cdot f}(O'_i)\}$ has the same joint distribution as the $n \times 1$ random vector with i^{th} component $t_{I_{\mathcal{O}_0} \cdot f}(O_i) - t_{I_{\mathcal{O}_0} \cdot f}(O'_i)$. Hence,

$$\begin{aligned} \mathbb{E}\left[\sup_{f \in \mathcal{S}_r} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i \left\{ t_{I_{\mathcal{O}_0} \cdot f}(O_i) - t_{I_{\mathcal{O}_0} \cdot f}(O'_i) \right\} \right| \right] &\leq 2\mathbb{E}\left\{ \sup_{f \in \mathcal{S}_r} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i t_{I_{\mathcal{O}_0} \cdot f}(O_i) \right| \right\} \\ &= 2\mathbb{E}\left\{ \sup_{f \in \mathcal{S}_r} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i t_f(O_i) \right| \right\} \end{aligned}$$

where $\mathcal{S}'_r := \{I_{\mathcal{O}_0} \cdot f : f \in \mathcal{S}_r\}$. Next, invoking the Ledoux-Talagrand contraction inequality (see inequality (5.61) of [Wainwright \(2019\)](#)), we obtain

$$\begin{aligned} 2\mathbb{E}\left\{ \sup_{f \in \mathcal{S}'_r} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i t_f(O_i) \right| \right\} &\leq 4L\mathbb{E}\left\{ \sup_{f \in \mathcal{S}'_r} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i f(O_i) \right| \right\} \\ &= 4L\mathbb{E}\left\{ \sup_{f \in \mathcal{S}_r} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i I_{\mathcal{O}_0}(O_i) \cdot f(O_i) \right| \right\} \\ &= 4L\mathbb{E}\left\{ \sup_{f \in \mathcal{F}_r} \left| \frac{1}{n} \sum_{i=1}^n \varepsilon_i I_{\mathcal{O}_0}(O_i) \cdot f(O_i) \right| \right\} = 4L\mathcal{R}_n(I_{\mathcal{O}_0} \cdot \mathcal{F}, r), \end{aligned}$$

where the second equality follows by applying Corollary [C.1](#) to the functions $t_f(o, \varepsilon) := \varepsilon f(o)$ and $\beta = 0$, noting that assumption 3 of the corollary holds because the map $f \mapsto f(o)$ is continuous (assumption 1.2). Thus, $\mathbb{E}\{Z_n(r)\} \leq 4L \cdot \mathcal{R}_n(I_{\mathcal{O}_0} \cdot \mathcal{F}, r)$. Therefore,

$$\mathbb{P}[Z_n(r) \geq 4L \cdot \mathcal{R}_n(I_{\mathcal{O}_0} \cdot \mathcal{F}, r) + u] \leq \mathbb{P}[Z_n(r) \geq \mathbb{E}\{Z_n(r)\} + u]$$

$$\begin{aligned}
&\leq 2 \exp \left(- \frac{nu^2}{8eL^2r^2 + 32ebL\mathbb{E}\{Z_n(r)\} + 8bLu} \right) \\
&\leq 2 \exp \left(- \frac{nu^2}{8eL^2r^2 + 128ebL^2\mathcal{R}_n(I_{\mathcal{O}_0} \cdot \mathcal{F}, r) + 8bLu} \right).
\end{aligned}$$

Because $\text{star}(I_{\mathcal{O}_0} \cdot \mathcal{F})$ is star-shaped, by the arguments in the proof of Lemma 13.6 of [Wainwright \(2019\)](#), the function $\delta \mapsto \frac{\mathcal{R}_n(\text{star}(I_{\mathcal{O}_0} \cdot \mathcal{F}), \delta)}{\delta}$ is non-decreasing on the interval $(0, \infty)$. Also, since $I_{\mathcal{O}_0} \cdot \mathcal{F} \subset \text{star}(I_{\mathcal{O}_0} \cdot \mathcal{F})$, then, for any $r \geq \delta_n$, it holds $\mathcal{R}_n(I_{\mathcal{O}_0} \cdot \mathcal{F}, r)/r \leq \mathcal{R}_n(\text{star}(I_{\mathcal{O}_0} \cdot \mathcal{F}), r)/r \leq \mathcal{R}_n(\text{star}(I_{\mathcal{O}_0} \cdot \mathcal{F}), \delta_n)/\delta_n \leq \frac{\delta_n}{b}$. Then, $\mathbb{P}\{Z_n(r) \geq \frac{4Lr\delta_n}{b} + u\} \leq \mathbb{P}\{Z_n(r) \geq 4L\mathcal{R}_n(I_{\mathcal{O}_0} \cdot \mathcal{F}, r) + u\} \leq 2 \exp \left(- \frac{nu^2}{8eL^2r^2 + 128ebL^2r\delta_n + 8bLu} \right)$. $\square$

Lemma C.8 (Lemma 12 of [Foster and Syrgkanis \(2023\)](#)). *Suppose the assumptions of Lemma C.7 hold. Let h_f and $Z_n(r)$ be defined as in Lemma C.7 and let δ_n be any solution to $\mathcal{R}_n(\text{star}(I_{\mathcal{O}_0} \cdot \mathcal{F}), \delta) \leq \frac{\delta^2}{b}$. For fixed $\theta > 1$ define the event $\mathcal{E} := \{\exists f \in \mathcal{F} : \|I_{\mathcal{O}_0} \cdot f\|_2 \geq \delta_n \text{ and } |\frac{1}{n} \sum_{i=1}^n h_f(O_i)| \geq \frac{5\theta L\delta_n}{b} \|I_{\mathcal{O}_0} \cdot f\|_2\}$. Also, set $c_1 = \{8b^2(1 + 17e)\}^{-1}$. Then, $P(\mathcal{E}) \leq 2\{1 + I\{\delta_n < b/\theta\} \lfloor \{\log(\theta)\}^{-1} \log(b/\delta_n) \rfloor\} \exp(-c_1 n \delta_n^2)$ where $\lfloor x \rfloor$ denotes the greatest integer smaller than or equal to x .*

Proof. This proof uses a peeling argument as in the proof of Lemma 12 of [Foster and Syrgkanis \(2023\)](#). Since $\|I_{\mathcal{O}_0} \cdot f\|_2 \leq \|f\|_2 \leq \|f\|_\infty$, Assumption 1.1 of Lemma C.7 implies that $\|I_{\mathcal{O}_0} \cdot f\|_2 \leq b$ for any $f \in \mathcal{F}$. Then, any $f \in \mathcal{F}$ satisfying $\|I_{\mathcal{O}_0} \cdot f\|_2 \geq \delta_n$ belongs to the set $\mathcal{S}_m := \{f \in \mathcal{F} : \theta^{m-1}\delta_n \leq \|I_{\mathcal{O}_0} \cdot f\|_2 \leq \theta^m\delta_n\}$ for some $m \in \{1, \dots, M\}$ where M is the smallest integer satisfying $b \leq \theta^M\delta_n$. We find such M as follows. If $\delta_n \geq b/\theta$ then $M = 1$ is such constant because $\theta > 1$. If $\delta_n < b/\theta$ then $b \leq \theta^M\delta_n$ iff $\log(b/\delta_n) \leq M \log(\theta)$ iff $\{\log(\theta)\}^{-1} \log(b/\delta_n) \leq M$. Therefore, we conclude that $M = 1 + I\{\delta_n < b/\theta\} \lfloor \{\log(\theta)\}^{-1} \log(b/\delta_n) \rfloor$. Note that M is a positive integer because $\theta > 1$ and $b/\delta_n > \theta$ when $\delta_n < b/\theta$.

Now, if the event $\mathcal{E} \cap \mathcal{S}_m$ occurs for some $m \in \{1, \dots, M\}$, then there exists f with

$\theta^{m-1}\delta_n \leq \|I_{\mathcal{O}_0} \cdot f\|_2 \leq r_m := \theta^m \delta_n$ such that $|\frac{1}{n} \sum_{i=1}^n h_f(O_i)| \geq \frac{5\theta L\delta_n}{b} \|I_{\mathcal{O}_0} \cdot f\|_2 \geq \frac{5\theta L\delta_n}{b} \frac{r_m}{\theta} = 5\frac{L\delta_n r_m}{b}$. Consequently, we have $P(\mathcal{E} \cap \mathcal{S}_m) \leq P[Z_n(r_m) \geq 5\frac{L\delta_n r_m}{b}]$. Applying Lemma C.7 with $r = r_m$ and $u = \frac{L\delta_n r_m}{b}$, and setting $c_1 = \{8b^2(1 + 17e)\}^{-1}$, we obtain that $\Pr(\mathcal{E} \cap \mathcal{S}_m)$ happens with probability at most $2 \exp\left\{-\frac{nL^2 r_m^2 \delta_n^2}{8b^2(eL^2 r_m^2 + 16eL^2 \delta_n r_m + L^2 \delta_n^2)}\right\} \leq 2 \exp(-c_1 n \delta_n^2)$ where the inequality holds since $\delta_n \leq r_m$. Then $P(\mathcal{E}) \leq \sum_{m=1}^M P(\mathcal{E} \cap \mathcal{S}_m) \leq 2M \exp(-c_1 n \delta_n^2) = 2\{1 + I\{\delta_n < b/\theta\} \lfloor \{\log(\theta)\}^{-1} \log(b/\delta_n) \rfloor\} \exp(-c_1 n \delta_n^2)$. $\square$

Lemma C.9 (Lemma 14 of Foster and Syrgkanis (2023)). *Suppose the assumptions of Lemma C.7 hold. Let h_f and $Z_n(r)$ be defined as in Lemma C.7 and let δ_n be any solution to the inequality $\mathcal{R}_n(\text{star}(I_{\mathcal{O}_0} \cdot \mathcal{F}), \delta) \leq \frac{\delta^2}{b}$. Set $c_1 = \{8b^2(1 + 17e)\}^{-1}$. Then,*

$$\left| \frac{1}{n} \sum_{i=1}^n h_f(O_i) \right| \leq \frac{6L\delta_n}{b} \{\|I_{\mathcal{O}_0} \cdot f\|_2 + \delta_n\} \quad \forall f \in \mathcal{F} \quad (\text{C.5})$$

happens with probability at least $1 - \{4 + 2I\{\delta_n < 5b/6\} \lfloor \{\log(6/5)\}^{-1} \log(b/\delta_n) \rfloor\} \exp(-c_1 n \delta_n^2)$.

Proof. Consider the events $\mathcal{E}_0 = \{Z_n(\delta_n) \geq 5L\delta_n^2/b\}$ and

$$\mathcal{E}_1 = \left\{ \exists f \in \mathcal{F} : \|I_{\mathcal{O}_0} \cdot f\|_2 \geq \delta_n \text{ and } \left| \frac{1}{n} \sum_{i=1}^n h_f(O_i) \right| \geq \frac{6L\delta_n}{b} \|I_{\mathcal{O}_0} \cdot f\|_2 \right\}.$$

If (C.5) is violated, then either $\mathcal{E}_0$ or $\mathcal{E}_1$ must occur. Applying Lemma C.7 with $r = \delta_n$ and $u = \frac{L\delta_n^2}{b}$ implies that $\mathcal{E}_0$ happens with probability at most $2 \exp(-c_1 n \delta_n^2)$ where $c_1 = \{8b^2(1 + 17e)\}^{-1}$. Applying Lemma C.8 with $\theta = 6/5$ we conclude that $\mathcal{E}_1$ happens with probability at most $\{2 + 2I\{\delta_n < 5b/6\} \lfloor \{\log(6/5)\}^{-1} \log(b/\delta_n) \rfloor\} \exp(-c_1 n \delta_n^2)$. Thus, (C.5) is violated with probability at most $\{4 + 2I\{\delta_n < 5b/6\} \lfloor \{\log(6/5)\}^{-1} \log(b/\delta_n) \rfloor\} \exp(-c_1 n \delta_n^2)$. This concludes the proof. $\square$

To state the next lemma, let B and B' be the constants defined in Assumption 4.5, and B_1 ,

b , and c_1 as in Theorem 4.4. Let $k_1 := \{\sup_{w \in \mathcal{W}} K_{\mathcal{H}}(w, w)\}^{1/2}$, $k_2 := \{\sup_{z \in \mathcal{Z}} K_{\mathcal{G}}(z, z)\}^{1/2}$, $c_0 := 9 \max\{\|s\|_{\infty}, \|r\|_{\infty} \cdot k_1 B, 2c^2 b\}/(c^2 b)$, and $c'_0 := c^2 \cdot c_0(2 + 3c_0 + 2 \max\{1, c_0/9\})$. The following lemma extends Lemma 10 of Bennett et al. (2023b) to allow for $\lambda_G \neq 0$.

Lemma C.10. *Under the assumptions of Theorem 4.4 and with the constants defined as in that Theorem, the following event happens for any $h \in \mathcal{H}_B$:*

$$\begin{aligned} & \frac{1}{4c^2} \left\{ \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 - \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 \right\} + \lambda_{\mathcal{H}} \left(\|\tilde{h}\|_{\mathcal{H}}^2 - \|h\|_{\mathcal{H}}^2 \right) \\ & \leq \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 + c_0 \delta_n \|I_0 \mathcal{T}(\tilde{h} - h)\|_2 + \lambda_G B_1^2 + c'_0 \delta_n^2. \end{aligned}$$

Proof. For any $h \in \mathcal{H}_B$, define $g_h := \frac{1}{2c^2} \mathcal{T}(h_0^\dagger - h)$. Since $\|h_0^\dagger - h\|_{\mathcal{H}} \leq \|h_0^\dagger\|_{\mathcal{H}} + \|h\|_{\mathcal{H}} \leq 2B$, Assumption 4.5 of the main text implies that $g_h \in \mathcal{G}_{B_1}$ where $B_1 := B'/c^2$. In addition, for any $h \in \mathcal{H}_B$, we have $\sup_{w \in \mathcal{W}} |h(w)| \leq k_1 B$, so the function class $\mathcal{H}_B$ is uniformly bounded by $k_1 B$. Similarly, the function class $\mathcal{G}_{B_1}$ is uniformly bounded by $k_2 B_1$.

By part 1 of Proposition 4.2 and Assumption 4.5, we have

$$\frac{1}{4c^2} \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 = \mathbb{E}_{P_0} \{ I_0(X) s(O) g_{\tilde{h}}(Z) - I_0(X) r(O) \tilde{h}(W) g_{\tilde{h}}(Z) - c^2 \cdot I_0(X) g_{\tilde{h}}(Z)^2 \},$$

where the expectation in the right hand-side is with respect to O , and $\tilde{h}$ is regarded as a fixed, non-random function.

Next, we apply Lemma C.9 to bound the expectation of each of the three terms in the right hand side of the last display. Specifically, with $c_1 := \{8b^2(1 + 17e)\}^{-1}$, $c_2 := 6L'/b$, $b := k_2 B_1$, $\xi(\delta_n) := \{4 + 2I\{\delta_n < 5b/6\} \lfloor \{\log(6/5)\}^{-1} \log(b/\delta_n) \rfloor\} \exp(-c_1 n \delta_n^2)$, and

$$L' = \max\{\|s\|_{\infty}, \|r\|_{\infty} \cdot k_1 B, 2c^2 \cdot b\}, \quad (\text{C.6})$$

applying Lemma C.9, we can conclude that

I. With probability at least $1 - \xi(\delta_n)$, the following event happens

$$\left| \mathbb{E}_n [I_0(X)s(O)g(Z) - \mathbb{E}_{P_0} \{I_0(X)s(O)g(Z)\}] \right| \leq c_2 \delta_n \{\|I_0 g\|_2 + \delta_n\} \quad \forall g \in \mathcal{G}_{B_1}. \quad (\text{C.7})$$

In particular, $\mathbb{E}_{P_0} \{I_0(X)s(O)g_{\tilde{h}}(Z)\} \leq \mathbb{E}_n \{I_0(X)s(O)g_{\tilde{h}}(Z)\} + c_2 \delta_n \{\|I_0 g_{\tilde{h}}\|_2 + \delta_n\}$ holds with probability at least $1 - \xi(\delta_n)$.

II. With probability at least $1 - \xi(\delta_n)$, for all $g \in \mathcal{G}_{B_1}$, the following event happens

$$\left| \mathbb{E}_n [I_0(X)r(O)\tilde{h}(W)g(Z) - \mathbb{E}_{P_0} \{I_0(X)r(O)\tilde{h}(W)g(Z)\}] \right| \leq c_2 \delta_n \{\|I_0 g\|_2 + \delta_n\}. \quad (\text{C.8})$$

Notably, $-\mathbb{E}_{P_0} \{I_0(X)r(O)\tilde{h}(W)g_{\tilde{h}}(Z)\} \leq -\mathbb{E}_n \{I_0(X)r(O)\tilde{h}(W)g_{\tilde{h}}(Z)\} + c_2 \delta_n \{\|I_0 g_{\tilde{h}}\|_2 + \delta_n\}$ holds with probability at least $1 - \xi(\delta_n)$.

III. With probability at least $1 - \xi(\delta_n)$, the following event happens

$$\left| \mathbb{E}_n [c^2 I_0(X)g(Z)^2 - \mathbb{E}_{P_0} \{c^2 I_0(X)g(Z)^2\}] \right| \leq c_2 \delta_n \{\|I_0 g\|_2 + \delta_n\} \quad \forall g \in \mathcal{G}_{B_1}. \quad (\text{C.9})$$

In particular, $-\mathbb{E}_{P_0} \{c^2 I_0(X)g_{\tilde{h}}(Z)^2\} \leq -\mathbb{E}_n \{c^2 I_0(X)g_{\tilde{h}}(Z)^2\} + c_2 \delta_n \{\|I_0 g_{\tilde{h}}\|_2 + \delta_n\}$ holds with probability at least $1 - \xi(\delta_n)$.

To apply Lemma C.9 to prove (I), (II), and (III), we define convenient function classes $\mathcal{F}$, maps $f \in \mathcal{L}^2(P_{0,O}) \mapsto t_f$, and sets $\mathcal{O}_0$. For these objects, we argue that the following conditions hold:

a. $\mathcal{F}$ is a separable metric space with metric denoted by d such that $\mathcal{F} \subset \mathcal{L}^2(P_{0,O})$ and:

- (a.1) $\mathcal{F}$ is uniformly bounded by b , (a.2) for each $o \in \mathcal{O}$, the evaluation functional $f \mapsto f(o)$ is continuous relative to the metric d , and (a.3) $f = 0$ belongs to $\mathcal{F}$.

- b. For each $f \in \mathcal{L}^2(P_{0,O})$, $t_f : \mathcal{O} \mapsto \mathbb{R}$ is a measurable function.
- c. For each $o \in \mathcal{O}$, the map $f \mapsto t_{I_{\mathcal{O}_0} \cdot f}(o)$ is continuous with respect to d at all $f \in \mathcal{F}$.
- d. The map $f \mapsto t_{I_{\mathcal{O}_0} \cdot f}$ is continuous from $(\mathcal{F}, d)$ to $\mathcal{L}^1(P_{0,O})$.
- e. t_f satisfies $t_f = 0$ when $f = 0$.
- f. $\exists L$ such that for any $f_1, f_2 \in \mathcal{L}^2(P_{0,O})$, $|t_{f_1}(o) - t_{f_2}(o)| \leq L |f_1(o) - f_2(o)| \forall o \in \mathcal{O}$.

(i) *Application of Lemma C.9 to show (I).* Define $\mathcal{F} := \mathcal{G}_{B_1}$, $t_f(o) := s(o) \cdot f(z)$, and $\mathcal{O}_0 := \{o \in \mathcal{O} : \text{the subvector } x \text{ of } o \text{ is in } \mathcal{X}_0\}$. We endow $\mathcal{F}$ with the metric d associated with the norm $\|\cdot\|_{\mathcal{G}}$ of the RKHS $\mathcal{G}$. To show that condition (a) holds first note that $\mathcal{F} = \mathcal{G}_{B_1} \subset \mathcal{L}^2(P_{0,Z}) \subset \mathcal{L}^2(P_{0,O})$, and $\mathcal{G}_{B_1}$ is separable because $\mathcal{G}$ is a separable RKHS by Assumption 4.6 (see Proposition 11.7 of Paulsen and Raghupathi (2016)). We already argued that (a.1) holds with $b = k_2 B_1$, (a.2) holds because the evaluation functionals are continuous in RKHSs, and (a.3) trivially holds.

Condition (b) is trivially satisfied. Condition (c) holds because $|t_{I_{\mathcal{O}_0} \cdot f_1}(o) - t_{I_{\mathcal{O}_0} \cdot f_2}(o)| = |I_{\mathcal{O}_0}(o)s(o)\{f_1(z) - f_2(z)\}| \leq \|s\|_{\infty}|f_1(z) - f_2(z)| \leq \|s\|_{\infty}k_2\|f_1 - f_2\|_{\mathcal{G}}$, where the last inequality holds because $\|f\|_{\infty} \leq k_2\|f\|_{\mathcal{F}} \forall f \in \mathcal{F}$. Condition (d) holds because $\mathbb{E}_{P_0}\{|t_{I_{\mathcal{O}_0} \cdot f_1}(O) - t_{I_{\mathcal{O}_0} \cdot f_2}(O)|\} \leq \sup_{o \in \mathcal{O}} |t_{I_{\mathcal{O}_0} \cdot f_1}(o) - t_{I_{\mathcal{O}_0} \cdot f_2}(o)| \leq \|s\|_{\infty} \cdot k_2 \cdot \|f_1 - f_2\|_{\mathcal{G}}$. Condition (e) trivially holds. Arguing as in the proof of (c), condition (f) holds with $L = \|s\|_{\infty}$.

Before applying Lemma C.9, note that $I_0 \cdot \mathcal{G}_{B_1}$ is star-shaped because $\mathcal{G}_{B_1}$ is star-shaped and multiplication by I_0 preserves this. Hence, its critical radius equals that of its star-shaped hull, so the lemma applies with an upper bound on the critical radius of $I_0 \cdot \mathcal{G}_{B_1}$.

Applying Lemma C.9, we obtain that the display (C.7) happens with probability at least $1 - \xi(\delta_n)$ with $c_2 = 6 \cdot \|s\|_{\infty}/b$ and $b = k_2 B_1$. Consequently, the display (C.7) also happens with probability at least $1 - \xi(\delta_n)$ with c_2 now equal to $6L'/b$ where L' is defined in (C.6).

(ii) *Application of Lemma C.9 to show (II).* Define $\mathcal{F}$, the metric d , and $\mathcal{O}_0$ as in (i).

Define $t_f(o) := r(o) \cdot \tilde{h}(w) \cdot f(z)$. Condition (a) holds as shown in (i). Condition (b) holds trivially. Conditions (c) and (d) hold as shown in (i), using in addition that $\sup_{w \in \mathcal{W}} |\tilde{h}(w)| \leq k_1 B$. Condition (e) is immediate. Arguing as in the proof of condition (c) in (i), (f) holds with $L = \|r\|_\infty \cdot k_1 B$. Applying Lemma C.9, we obtain that the display (C.8) happens with probability at least $1 - \xi(\delta_n)$ with $c_2 = 6\|r\|_\infty k_1 B/b$ and $b = k_2 B_2$. Then, (C.8) also happens with probability at least $1 - \xi(\delta_n)$ with c_2 now equal to $6L'/b$ where L' is defined in (C.6).

(iii) *Application of Lemma C.9 to show (III).* Define $\mathcal{F}$, d , and $\mathcal{O}_0$ as in (i). Define $t_f(o) := c^2 \cdot f(z)^2$. Condition (a) holds as shown in (i), (b) holds trivially, (c) holds because $|t_{I_{\mathcal{O}_0} \cdot f_1}(o) - t_{I_{\mathcal{O}_0} \cdot f_2}(o)| = |c^2 I_{\mathcal{O}_0}(o) \{f_1(z)^2 - f_2(z)^2\}| \leq c^2 |f_1(z) + f_2(z)| \cdot |f_1(z) - f_2(z)| \leq 2c^2 k_2 B_1 |f_1(z) - f_2(z)| \leq 2c^2 k_2^2 B_1 \|f_1 - f_2\|_{\mathcal{G}}$, and (d) holds because $\mathbb{E}_{P_0}\{|t_{I_{\mathcal{O}_0} \cdot f_1}(O) - t_{I_{\mathcal{O}_0} \cdot f_2}(O)|\} \leq \sup_{o \in \mathcal{O}} |t_{I_{\mathcal{O}_0} \cdot f_1}(o) - t_{I_{\mathcal{O}_0} \cdot f_2}(o)| \leq 2c^2 \cdot k_2^2 \cdot B_1 \cdot \|f_1 - f_2\|_{\mathcal{G}}$. Condition (e) is trivial and condition (f) holds with $L = 2c^2 \cdot b$ arguing as in the proof of (c).

Applying Lemma C.9, we obtain that the display (C.9) happens with probability at least $1 - \xi(\delta_n)$ with $c_2 = 6 \cdot 2c^2 \cdot b/b$ and $b = k_2 B_1$. Consequently, the display (C.9) also happens with probability at least $1 - \xi(\delta_n)$ with c_2 now equal to $6L'/b$ where L' is defined in (C.6).

Setting $c_3 = 3c_2 = 18L'/b$ with L' defined in (C.6), with probability at least $1 - 3\xi(\delta_n)$, we have $\frac{1}{4c^2} \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 \leq \mathbb{E}_n\{I_0(X)s(O)g_{\tilde{h}}(Z) - I_0(X)r(O)\tilde{h}(W)g_{\tilde{h}}(Z) - c^2 I_0(X)g_{\tilde{h}}(Z)^2\} + c_3 \delta_n \|I_0 g_{\tilde{h}}\|_2 + c_3 \delta_n^2$.

Then, the following event happens with probability at least $1 - 3\xi(\delta_n)$, for any $h \in \mathcal{H}_B$:

$$\begin{aligned} \frac{1}{4c^2} \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 &\leq \mathbb{E}_n\{I_0(X)s(O)g_{\tilde{h}}(Z) - I_0(X)r(O)\tilde{h}(W)g_{\tilde{h}}(Z) - c^2 I_0(X)g_{\tilde{h}}(Z)^2\} \\ &\quad - \lambda_{\mathcal{G}} \|g_{\tilde{h}}\|_{\mathcal{G}}^2 + \lambda_{\mathcal{G}} \|g_{\tilde{h}}\|_{\mathcal{G}}^2 + c_3 (\delta_n \|I_0 \cdot g_{\tilde{h}}\|_2 + \delta_n^2) \\ &\leq \sup_{g \in \mathcal{G}} [\mathbb{E}_n\{I_0(X)s(O)g(Z) - I_0(X)r(O)\tilde{h}(W)g(Z) - c^2 I_0(X)g(Z)^2\} - \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2] \\ &\quad + \lambda_{\mathcal{G}} B_1^2 + c_3 (\delta_n \|I_0 \cdot g_{\tilde{h}}\|_2 + \delta_n^2) \end{aligned}$$

$$\begin{aligned}
&\leq \sup_{g \in \mathcal{G}} [\mathbb{E}_n \{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - c^2 I_0(X)g(Z)^2\} - \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2] \\
&\quad + \lambda_{\mathcal{G}} B_1^2 + \lambda_{\mathcal{H}} (\|h\|_{\mathcal{H}}^2 - \|\tilde{h}\|_{\mathcal{H}}^2) + c_3 (\delta_n \|I_0 \cdot g_{\tilde{h}}\|_2 + \delta_n^2). \tag{C.10}
\end{aligned}$$

Note that the second inequality is always true because $g_{\tilde{h}} \in \mathcal{G}_{B_1}$, and the third inequality is also true because $\tilde{h}$ is the empirical risk minimizer defined in (4.12).

Set $c_4 := \max\{c_3[1 + c_3/(2c^2)], 5c_2^2/c^2\}$. For any $h \in \mathcal{H}_B$, we define

$$T_n(h) := \sup_{g \in \mathcal{G}} \mathbb{E}_{P_0} \{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - \frac{c^2}{2} \cdot I_0(X)g(Z)^2\} + c_4 \delta_n^2,$$

$$A_n(h) := \sup_{g \in \mathcal{G}_{B_1}} [\mathbb{E}_n \{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - c^2 I_0(X)g(Z)^2\} - \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2]$$

$$B_n(h) := \sup_{g \in \mathcal{G} \setminus \mathcal{G}_{B_1}} [\mathbb{E}_n \{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - c^2 I_0(X)g(Z)^2\} - \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2],$$

and

$$C_n(h) := \sup_{g \in \mathcal{G}} [\mathbb{E}_n \{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - c^2 \cdot I_0(X)g(Z)^2\} - \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2].$$

Below we will show that, with probability at least $1 - 3\xi(\delta_n)$, each of the following event happens: $A_n(h) \leq T_n(h)$ and $B_n(h) \leq T_n(h)$. This will then imply that with probability at least $1 - 6\xi(\delta_n)$ for any $h \in \mathcal{H}_B$:

$$C_n(h) \leq T_n(h). \tag{C.11}$$

Assume for now that the last assertion holds. By part 1 of Proposition 4.2 and Assumption 4.5, we have that $\sup_{g \in \mathcal{G}} \mathbb{E}_{P_0} \{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - \frac{c^2}{2} I_0(X)g(Z)^2\} = \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2$. Combining this equality with the inequalities in (C.10) and (C.11), we

then conclude that with probability at least $1 - 9\xi(\delta_n)$, for any $h \in \mathcal{H}_B$:

$$\begin{aligned} \frac{1}{4c^2} \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 &\leq \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 + \lambda_{\mathcal{H}} \left(\|h\|_{\mathcal{H}}^2 - \|\tilde{h}\|_{\mathcal{H}}^2 \right) + \lambda_{\mathcal{G}} B_1^2 \\ &\quad + c_3 \cdot \delta_n \|I_0 \cdot g_{\tilde{h}}\|_2 + (c_3 + c_4) \delta_n^2. \end{aligned} \quad (\text{C.12})$$

Furthermore, by the triangle inequality $\|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2 \leq \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2 + \|I_0 \mathcal{T}(\tilde{h} - h)\|_2$ and since $\|I_0 \cdot g_{\tilde{h}}\|_2 = \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2$, the inequality (C.12) implies that the following event happens with probability at least $1 - 9\xi(\delta_n)$ for any $h \in \mathcal{H}_B$

$$\begin{aligned} \frac{1}{4c^2} \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 &\leq \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 + \lambda_{\mathcal{H}} \left(\|h\|_{\mathcal{H}}^2 - \|\tilde{h}\|_{\mathcal{H}}^2 \right) + \lambda_{\mathcal{G}} B_1^2 \\ &\quad + \frac{c_3}{2c^2} \left\{ \delta_n \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2 + \delta_n \|I_0 \mathcal{T}(\tilde{h} - h)\|_2 \right\} + (c_3 + c_4) \delta_n^2. \end{aligned}$$

By the AM-GM inequality $\frac{c_3}{2c^2} \delta_n \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2 \leq \frac{1}{4c^2} \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 + \frac{c_3^2}{4c^2} \delta_n^2$. Hence, the following event happens with probability at least $1 - 9\xi(\delta_n)$ for all $h \in \mathcal{H}_B$:

$$\begin{aligned} \frac{1}{4c^2} \left\{ \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 - \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 \right\} &+ \lambda_{\mathcal{H}} \left(\|\tilde{h}\|_{\mathcal{H}}^2 - \|h\|_{\mathcal{H}}^2 \right) \\ &\leq \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 + \frac{c_3}{2c^2} \delta_n \|I_0 \mathcal{T}(\tilde{h} - h)\|_2 + \lambda_{\mathcal{G}} B_1^2 + \left(c_3 + c_4 + \frac{c_3^2}{4c^2} \right) \delta_n^2, \end{aligned}$$

thus, arriving at the desired inequality.

We now turn to prove that $A_n(h) \leq T_n(h)$ holds with probability at least $1 - 3\xi(\delta_n)$. Fix $h \in \mathcal{H}_B$ and take $g \in \mathcal{G}_{B_1}$. Applying Lemma C.9 as we did to show (I), (II), and (III), and noting that $-\lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2 \leq 0$, we have that the following event happens with probability at least $1 - 3\xi(\delta_n)$,

$$\mathbb{E}_n \{ I_0(X) s(O) g(Z) - I_0(X) r(O) h(W) g(Z) - c^2 I_0(X) g(Z)^2 \} - \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2 \quad (\text{C.13})$$

$$\leq E_{P_0} \{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - c^2 I_0(X)g(Z)^2\} + c_3 \delta_n \|I_0 \cdot g\|_2 + c_3 \delta_n^2.$$

By the AM-GM inequality, $c_3 \delta_n \|I_0 \cdot g\|_2 \leq \frac{c^2}{2} \|I_0 \cdot g\|_2^2 + \frac{c_3^2}{2c^2} \delta_n^2$. Hence, the right-hand side of (C.13) is upper bounded by

$$\sup_{g \in \mathcal{G}} \mathbb{E}_{P_0} \left\{ I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - \frac{c^2}{2} I_0(X)g(Z)^2 \right\} + \left(c_3 + \frac{c_3^2}{2c^2} \right) \delta_n^2 \leq T_n(h).$$

Since $g \in \mathcal{G}_{B_1}$ was arbitrary, $A_n(h) \leq T_n(h)$ happens with probability at least $1 - 3\xi(\delta_n)$.

Turn now to prove that $B_n(h) \leq T_n(h)$ holds with probability at least $1 - 3\xi(\delta_n)$. Fix $h \in \mathcal{H}_B$ and take $g \in \mathcal{G} \setminus \mathcal{G}_{B_1}$. Then, $\|g\|_{\mathcal{G}} > B_1$. Define $\tilde{g} := B_1 \cdot g / \|g\|_{\mathcal{G}}$. Since $\mathcal{G}$ is star-shaped and $B_1 / \|g\|_{\mathcal{G}} \in (0, 1)$, $\tilde{g} \in \mathcal{G}$. Furthermore, $\|\tilde{g}\|_{\mathcal{G}} = B_1$, and so $\tilde{g} \in \mathcal{G}_{B_1}$. Applying Lemma C.9 as we did to prove (I) and (II), we obtain that, with probability at least $1 - 2\xi(\delta_n)$, the following happens:

$$\begin{aligned} & \mathbb{E}_n \{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z)\} \\ &= \frac{\|g\|_{\mathcal{G}}}{B_1} \cdot \mathbb{E}_n \{I_0(X)s(O)\tilde{g}(Z) - I_0(X)r(O)h(W)\tilde{g}(Z)\} \\ &\leq \frac{\|g\|_{\mathcal{G}}}{B_1} \cdot [\mathbb{E}_{P_0} \{I_0(X)s(O)\tilde{g}(Z) - I_0(X)r(O)h(W)\tilde{g}(Z)\} + 2c_2 \delta_n \|I_0 \cdot \tilde{g}\|_2 + 2c_2 \delta_n^2] \\ &= \mathbb{E}_{P_0} \{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z)\} + 2c_2 \delta_n \|I_0 \cdot g\|_2 + \frac{2c_2}{B_1} \|g\|_{\mathcal{G}} \cdot \delta_n^2, \end{aligned}$$

where $c_2 = 6L'/b$, with L' defined in (C.6). Applying Lemma C.9 as we did to prove (III), with probability at least $1 - \xi(\delta_n)$, the following event happens:

$$\begin{aligned} -\mathbb{E}_n \{c^2 I_0(X)g(Z)^2\} &= -\frac{\|g\|_{\mathcal{G}}^2}{B_1^2} \mathbb{E}_n \{c^2 I_0(X)\tilde{g}(Z)^2\} \\ &\leq \frac{\|g\|_{\mathcal{G}}^2}{B_1^2} [-\mathbb{E}_{P_0} \{c^2 I_0(X)\tilde{g}(Z)^2\} + 12c^2 \delta_n (\|I_0 \cdot \tilde{g}\|_2 + \delta_n)] \end{aligned}$$

$$= -\mathbb{E}_{P_0}\{c^2 I_0(X)g(Z)^2\} + \frac{12c^2}{B_1}\delta_n\|g\|_{\mathcal{G}} \cdot \|I_0 \cdot g\|_2 + \frac{12c^2}{B_1^2}\|g\|_{\mathcal{G}}^2 \cdot \delta_n^2.$$

Hence, with probability at least $1 - 3\xi(\delta_n)$, the following event happens:

$$\begin{aligned} & \mathbb{E}_n\{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - c^2 I_0(X)g(Z)^2\} - \lambda_{\mathcal{G}}\|g\|_{\mathcal{G}}^2 \\ & \leq \mathbb{E}_P\{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - c^2 I_0(X)g(Z)^2\} - \lambda_{\mathcal{G}}\|g\|_{\mathcal{G}}^2 \quad (\text{C.14}) \\ & \quad + 2c_2\delta_n\|I_0 \cdot g\|_2 + \frac{2c_2}{B_1}\|g\|_{\mathcal{G}} \cdot \delta_n^2 + \frac{12c^2}{B_1}\delta_n\|g\|_{\mathcal{G}} \cdot \|I_0 \cdot g\|_2 + \frac{12c^2}{B_1^2}\|g\|_{\mathcal{G}}^2 \cdot \delta_n^2. \end{aligned}$$

Applying the AM-GM inequality, we obtain

$$\begin{aligned} 2c_2\delta_n\|I_0 \cdot g\|_2 & \leq \frac{c^2}{4}\|I_0 \cdot g\|_2^2 + \frac{4c_2^2}{c^2}\delta_n^2, \\ \frac{2c_2}{B_1}\|g\|_{\mathcal{G}} \cdot \delta_n^2 & \leq \frac{c_2^2}{c^2}\delta_n^2 + \frac{c^2}{B_1^2}\|g\|_{\mathcal{G}}^2 \cdot \delta_n^2, \quad \text{and} \\ \frac{12c^2}{B_1}\delta_n\|g\|_{\mathcal{G}} \cdot \|I_0 \cdot g\|_2 & \leq \frac{c^2}{4}\|I_0 \cdot g\|_2^2 + \frac{144c^2}{B_1^2}\|g\|_{\mathcal{G}}^2 \cdot \delta_n^2. \end{aligned}$$

Using these inequalities and the fact that $\lambda_{\mathcal{G}} \geq 157c^2 \cdot \delta_n^2 / B_1^2$, the right-hand side of (C.14) is upper bounded by

$$\begin{aligned} & E_{P_0}\{I_0(X)s(O)g(Z) - I_0(X)r(O)h(W)g(Z) - \frac{c^2}{2}I_0(X)g(Z)^2\} + \frac{5c_2^2}{c^2}\delta_n^2 \\ & \leq \sup_{g \in \mathcal{G}} E_{P_0}\{m(O; g) - I_0(X)r(O)h(W)g(Z) - \frac{c^2}{2}I_0(X)g(Z)^2\} + \frac{5c_2^2}{c^2}\delta_n^2 \leq T_n(h). \end{aligned}$$

Thus, $B_n(h) \leq T_n(h)$ holds with probability at least $1 - 3\xi(\delta_n)$. This completes the proof. $\square$

We now state a more general version of Lemma C.10 that applies for an arbitrary measurable map $m(o; g)$, linear in g , that satisfies Assumption 4.7.

Lemma C.11. *Under the assumptions of Theorem 4.5 and with $\lambda_{\mathcal{G}} \geq 157c^2 / B_1^2$, with proba-*

bility at least $1 - \{36 + 18 \cdot I\{\delta_n < 5b/6\} \lfloor \log(6/5) \rfloor^{-1} \log(b/\delta_n) \rfloor\} \exp(-c_1 n \delta_n^2)$, the following event happens for any $h \in \mathcal{H}_B$:

$$\begin{aligned} \frac{1}{4c^2} \left\{ \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 - \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 \right\} + \lambda_{\mathcal{H}} \left(\|\tilde{h}\|_{\mathcal{H}}^2 - \|h\|_{\mathcal{H}}^2 \right) \\ \leq \frac{1}{2c^2} \|I_0 \mathcal{T}(h_0^\dagger - h)\|_2^2 + c_0 \delta_n \|I_0 \mathcal{T}(\tilde{h} - h)\|_2 + \lambda_{\mathcal{G}} B_1^2 + c'_0 \delta_n^2, \end{aligned}$$

where $\tilde{h} = \arg \min_{h \in \mathcal{H}_B} \max_{g \in \mathcal{G}} \{\mathbb{E}_n[\{m(O; g) - I_0(X)r(O)h(W)g(Z) - c^2 I_0(X)g(Z)^2\}] - \lambda_{\mathcal{G}} \|g\|_{\mathcal{G}}^2 + \lambda_{\mathcal{H}} \|h\|_{\mathcal{H}}^2\}$, $c_0 := 3 \cdot \max\{3, B_2 + 2\} \cdot \max\{1 \|r\|_{\infty} \cdot k_1 B, 2c^2 b\} / (c^2 b)$, $b = B_1 \max\{k_2, B_3\}$, with B_2 and B_3 the constants specified in Assumption 4.7, and c_1 and c'_0 are defined as in Lemma C.10.

Proof. The proof proceeds similarly to that of Lemma C.10, with minor changes to handle the terms $\mathbb{E}_{P_0}\{m(O; g)\}$ and $\mathbb{E}_n\{m(O; g)\}$. Here, $\xi(\delta_n)$ is as defined in C.10, and the constants c_1 and c_2 are as in Lemma C.10, except that L' and b are now $L' := \max\{1, \|r\|_{\infty} \cdot k_1 B, 2c^2 \cdot b\}$ and $b := B_1 \max\{k_2, B_3\}$, with B_3 being the constant specified in Assumption 4.7.

To use Lemma C.9 to bound $\mathbb{E}_{P_0}\{m(O; g_{\tilde{h}})\}$, define $\mathcal{F} := \{f(\cdot) := m(\cdot; g) : \mathcal{O} \mapsto \mathbb{R} \mid g \in \mathcal{G}_{B_1}\}$, $t_f := f$, and $\mathcal{O}_0 := \mathcal{O}$ (equivalently $I_{\mathcal{O}_0} = 1$). We verify conditions (a)–(f) listed in the proof of Lemma C.10, which are required for Lemma C.9. Condition (a) follows immediately from part 2 of Assumption 4.7. Condition (a.3) holds because $g = 0$ belongs to $\mathcal{G}_{B_1}$ and therefore $f_0 := 0 = m(\cdot, 0)$ belongs to $\mathcal{F}$ by the linearity of m . Next, for any $f \in \mathcal{F}$, $f = m(\cdot, g_f)$ for some $g \in \mathcal{G}_{B_1}$. Then,

$$|f(o)| = |f(o) - f_0(o)| \leq B_3 \cdot d(f, f_0) = B_3 \cdot d(m(o; g_f), m(o, 0)) \leq B_3 \|g_f - 0\|_{\mathcal{G}} \leq B_3 B_1,$$

where the first and second inequalities follow by part 2 of Assumption 4.7, and the third holds because $g_f \in \mathcal{G}_{B_1}$. Thus, $\mathcal{F}$ is uniformly bounded by $B_3 B_1 \leq b$, and so it satisfies condition

(a.1). Condition (a.2) follows from part 2 of Assumption 4.7, which implies for all $o \in \mathcal{O}$ and $f_1, f_2 \in \mathcal{F}$ that $|f_1(o) - f_2(o)| \leq \sup_{o \in \mathcal{O}} |f_1(o) - f_2(o)| \leq B_3 \cdot d(f_1, f_2)$. Condition (b) holds by the assumption on m specified in Section 4.2. Condition (c) holds because $t_{I_{\mathcal{O}_0} \cdot f} = f$ and therefore (c) is the same as (a.2) that we have already argued to hold. Condition (d) holds because the last display implies $\mathbb{E}_{P_0}\{|f_1(O) - f_2(O)|\} \leq \sup_{o \in \mathcal{O}} |f_1(o) - f_2(o)| \leq B_3 \cdot d(f_1, f_2)$. Condition (e) trivially holds because $t_f = f$. Likewise, (f) trivially holds with $L = 1 \leq L'$.

By Lemma C.9, $\mathbb{E}_{P_0}\{m(O; g_{\tilde{h}})\} \leq \mathbb{E}_n\{m(O; g_{\tilde{h}})\} + c_2 \delta_n \{\|m(O; g_{\tilde{h}})\|_2 + \delta_n\}$ happens with probability at least $1 - \xi(\delta_n)$. Also, $\|m(O; g_{\tilde{h}})\|_2 \leq B_2 \|I_0 \cdot \tilde{g}\|_2$ by part 1 of Assumption 4.7. Then, $\mathbb{E}_{P_0}\{m(O; g_{\tilde{h}})\} \leq \mathbb{E}_n\{m(O; g_{\tilde{h}})\} + B_2 c_2 \cdot \delta_n \|I_0 \cdot g_{\tilde{h}}\|_2 + c_2 \delta_n^2$ holds with probability at least $1 - \xi(\delta_n)$. Combining the latter with the events in (C.8) and (C.9) from the proof of Lemma C.10, each of which holds with probability at least $1 - \xi(\delta_n)$, we conclude that, with probability at least $1 - 3\xi(\delta_n)$,

$$\begin{aligned} \frac{1}{4c^2} \|I_0 \mathcal{T}(h_0^\dagger - \tilde{h})\|_2^2 &\leq \mathbb{E}_n\{I_0(X)s(O)g_{\tilde{h}}(Z) - I_0(X)r(O)\tilde{h}(W)g_{\tilde{h}}(Z) - c^2 \cdot I_0(X)g_{\tilde{h}}(Z)^2\} \\ &\quad + c_2 \delta_n \|I_0 \cdot g_{\tilde{h}}\|_2 + c_2 \delta_n^2, \end{aligned}$$

where $c_3 := \max\{B_2 + 2, 3\} \cdot c_2$. The remainder of the proof then follows exactly as in Lemma C.10, with this updated constant c_3 replacing the value $c_3 = 3c_2 = 18L'/b$ used there. $\square$

C.3 Summary of Existing Regularized Adversarial Estimation Proposals

C.3.1 Specifications of Existing Adversarial Estimation Approaches

In the following Table C.1 we summarize existing adversarial proposals. The columns $\mathcal{H}$ and $\mathcal{G}$ indicate the estimation class $\mathcal{H}$ and the test class $\mathcal{G}$ used by each proposal. The symbol $+$ indicates that in addition to the listed classes, the proposals were made for additional function

classes. The column $RMSE$ provides the derived finite-sample bounds for the optimal choices of the tuning parameters in terms of the critical radius, the coefficient β of the assumed source condition, if any, or the measure of ill-posedness $\tau(\delta_n)$ as a function of δ_n .

The column entitled $\lambda_{\mathcal{H}} > 0$ *analysis/impl.?* indicates whether the estimator for which the theoretical analysis was presented penalizes the estimated function h , i.e., whether $\lambda_{\mathcal{H}} > 0$, and whether the estimation proposal for implementation penalizes as well, and the norm used in the penalization. The column $\lambda_{\mathcal{G}} > 0$ *analysis/impl.?* provides the same information but this time for the test function class $\mathcal{G}$. The last column indicates whether numerical experiments were provided in the paper.

Table C.1: Existing Adversarial Estimation Approaches

Reference	$\mathcal{H}$	$\mathcal{G}$	RMSE	$\lambda_{\mathcal{H}} > 0$ analysis/impl.?	$\lambda_{\mathcal{G}} > 0$ analysis/impl.?	Sims
Regularized Adversarial Stabilized Estimators						
Bennett et al. (2019)	NNs	NNs	$o_p(1)$	No/No	No/No	Yes
Muandet et al. (2020b)	RKHS	RKHS	$o_p(1)$	No/Yes ($\ \cdot\ _{\mathcal{H}}$)	No/Yes ($\ \cdot\ _{\mathcal{G}}$)	Yes
Dikkala et al. (2020)	RKHS+	RKHS+	$\tau(\delta_n)$	Yes/Yes ($\ \cdot\ _{\mathcal{H}}$)	Yes/Yes ($\ \cdot\ _{\mathcal{G}}$)	Yes
Liao et al. (2020)	NNs	NNs	$O_p\left(\frac{\min\{\beta, 2\}}{\delta^{2(1+\min\{\beta, 2\})}}\right)$	Yes/Yes ($\ \cdot\ _2$)	No/No	No
Kallus et al. (2021)	NNs+	RKHS+	$O_p(\tau^* \cdot \delta_n)$	No/Yes ($\ \cdot\ _{\mathcal{H}}$)	No/Yes ($\ \cdot\ _{\mathcal{G}}$)	Yes
Bennett et al. (2022)	NNs+	NNs+	$O_p(\delta_n^{1/2})$	Yes/Yes ($\ \cdot\ _{\mathcal{H}}$)	Yes/Yes ($\ \cdot\ _{\mathcal{G}}$)	Yes
Ghassami et al. (2022)	RKHS	RKHS	$\tau(\delta_n)$	Yes/Yes ($\ \cdot\ _{\mathcal{H}}$)	Yes/Yes ($\ \cdot\ _{\mathcal{G}}$)	Yes
Bennett et al. (2023a)	RKHS+	RKHS+	$O_p(n^{-1/4})$	Yes/Yes ($\ \cdot\ _2$)	No/No	No
Bennett et al. (2023b)	RKHS+	RKHS+	$O_p\left(\frac{\min\{\beta, 2t\}}{\delta_n^{2+\min\{\beta, 2t\}}}\right)$	Yes/Yes ($\ \cdot\ _2$)	No/Yes ($\ \cdot\ _{\mathcal{G}}$)	No
This work	RKHS	RKHS	$O_p\left(\frac{\min\{\beta, 1\}}{\delta_n^{1+\min\{\beta, 1\}}}\right)$	Yes/Yes ($\ \cdot\ _{\mathcal{H}}$)	Yes/Yes ($\ \cdot\ _{\mathcal{G}}$)	Yes*
Adversarial Maximal Moment Restriction Approaches						
Zhang et al. (2023)	RKHS+	RKHS+	Unclear	Yes/Yes ($\ \cdot\ _{\mathcal{H}}$)	No/No	Yes
Mastouri et al. (2021)	RKHS	RKHS	$O_p\left(n^{-\min\{\frac{\beta}{2\beta+4}, \frac{1}{4}\}}\right)$	Yes/Yes ($\ \cdot\ _{\mathcal{H}}$)	No/No	Yes
Park et al. (2024)	RKHS	RKHS	$O_p\left(n^{-\frac{\min\{\beta, 2\}}{2(1+\min\{\beta, 2\})}}\right)$	Yes/Yes ($\ \cdot\ _{\mathcal{H}}$)	No/No	Yes

*Numerical experiments provided in Chapter 3.

C.3.2 Comparison of Assumptions Across Regularized Adversarial Stabilized Estimation Approaches

Table C.2 below summarizes the key assumptions made in the listing papers for deriving finite-sample bounds. The column *Realizability* refers to whether or not it is assumed that the target solution to the ill-posedness problem is assumed to lie in the estimation class $\mathcal{H}$. In the row corresponding to the paper of [Bennett et al. \(2023b\)](#), the caption indicates that these authors assume that the solutions to the surrogate regularized problem (4.4) is in $\mathcal{H}$ for all $\lambda_{\mathcal{H}}$. These authors did not assume that the solution to the ill-posed problem belongs to $\mathcal{H}$. The columns *Uniqueness*, *Closedness*, and *Boundedness* are clear. The column $\mathcal{H}_B$ indicates that the outer minimization is conducted over the set of functions in $\mathcal{H}$ with norm of the class bounded by B . Analogously, the column $\mathcal{G}_{B'}$ refers to the test function class in the inner maximization. The last column summarizes the key assumption that enable derivation of RMSE bounds.

Table C.2: Comparison of Assumptions for Finite-Sample Bounds

Reference	Realizability	Uniqueness	Closedness	Boundedness	$\mathcal{H}_B$	$\mathcal{G}_{B'}$	Condition Enabling RMSE Bounds
Dikkala et al. (2020)	Yes	Yes	Global	Lipschitz	Yes	No	Measure of ill-posedness
Liao et al. (2020)	Yes	Yes	Local	Yes	Yes	Yes	Source-condition
Kallus et al. (2021)	Yes	Yes	Local	Yes	Yes	Yes	Measure of ill-posedness
Bennett et al. (2022)	Yes	No	Local	Yes	Yes	Yes	Source-condition ($\beta = 1$)
Ghassami et al. (2022)	Yes	Yes	Global	Lipschitz	Yes	No	Measure of ill-posedness
Bennett et al. (2023b)	For all $h_{\lambda_{\mathcal{H}}}$	No	Local	Yes	Yes	Yes	Source-condition
This work	Yes	No	Local	Yes	Yes	No	Source-condition

BIBLIOGRAPHY

- Ai, C. and Chen, X. (2003). Efficient estimation of models with conditional moment restrictions containing unknown functions. *Econometrica*, 71(6):1795–1843.
- Ai, C. and Chen, X. (2007). Estimation of possibly misspecified semiparametric conditional moment restriction models with different conditioning variables. *Journal of Econometrics*, 141(1):5–43.
- Baeten, J. M., Donnell, D., Ndase, P., Mugo, N. R., Campbell, J. D., Wangisi, J., Tappero, J. W., Bukusi, E. A., Cohen, C. R., Katabira, E., et al. (2012). Antiretroviral prophylaxis for hiv prevention in heterosexual men and women. *New England Journal of Medicine*, 367(5):399–410.
- Bartlett, P. L., Bousquet, O., and Mendelson, S. (2005). Local rademacher complexities. *The Annals of Statistics*.
- Bekker, L.-G., Das, M., Abdool Karim, Q., Ahmed, K., Batting, J., Brumskine, W., Gill, K., Harkoo, I., Jaggernath, M., Kigozi, G., et al. (2024). Twice-yearly lenacapavir or daily f/taf for hiv prevention in cisgender women. *New England Journal of Medicine*, 391(13):1179–1192.
- Bennett, A., Kallus, N., Mao, X., Newey, W., Syrgkanis, V., and Uehara, M. (2022). Inference on strongly identified functionals of weakly identified functions. *arXiv preprint arXiv:2208.08291*.

- Bennett, A., Kallus, N., Mao, X., Newey, W., Syrgkanis, V., and Uehara, M. (2023a). Minimax instrumental variable regression and l_2 convergence guarantees without identification or closedness. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 2291–2318. PMLR.
- Bennett, A., Kallus, N., Mao, X., Newey, W., Syrgkanis, V., and Uehara, M. (2023b). Source condition double robust inference on functionals of inverse problems. *arXiv preprint arXiv:2307.13793*.
- Bennett, A., Kallus, N., and Schnabel, T. (2019). Deep generalized method of moments for instrumental variable analysis. *Advances in neural information processing systems*, 32.
- Borwein, J. and Lewis, A. (2006). *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Springer.
- Boyd, S. (2004). Convex optimization. *Cambridge UP*.
- Brittain, E. H., Fay, M. P., and Follmann, D. A. (2012). A valid formulation of the analysis of noninferiority trials under random effects meta-analysis. *Biostatistics*, 13(4):637–649.
- Carrasco, M., Florens, J.-P., and Renault, E. (2007). Linear inverse problems in structural econometrics estimation based on spectral decomposition and regularization. *Handbook of econometrics*, 6:5633–5751.
- Cavalier, L. (2011). Inverse problems in statistics. In *Inverse Problems and High-Dimensional Estimation: Stats in the Château Summer School, August 31-September 4, 2009*, pages 3–96. Springer.
- Chen, X. and Pouzo, D. (2012). Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals. *Econometrica*, 80(1):277–321.

- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters.
- Chernozhukov, V., Newey, W., Quintas-Martinez, V. M., and Syrgkanis, V. (2022a). Riesznet and forestriesz: Automatic debiased machine learning with neural nets and random forests. In *International Conference on Machine Learning*, pages 3901–3914. PMLR.
- Chernozhukov, V., Newey, W. K., and Singh, R. (2022b). Automatic debiased machine learning of causal and structural effects. *Econometrica*, 90(3):967–1027.
- Chernozhukov, V., Newey, W. K., and Singh, R. (2022c). Debiased machine learning of global and local parameters using regularized riesz representers. *The Econometrics Journal*, 25(3):576–601.
- Chernozhukov, V., Newey, W. K., and Singh, R. (2023). A simple and general debiased machine learning theorem with finite-sample guarantees. *Biometrika*, 110(1):257–264.
- Conway, J. B. (2012). *A course in abstract analysis*, volume 141. American Mathematical Soc.
- Corey, L., Gilbert, P. B., Juraska, M., Montefiori, D. C., Morris, L., Karuna, S. T., Edupuganti, S., Mgodhi, N. M., DeCamp, A. C., Rudnicki, E., et al. (2021). Two randomized trials of neutralizing antibodies to prevent hiv-1 acquisition. *New England Journal of Medicine*, 384(11):1003–1014.
- Cui, Y., Pu, H., Shi, X., Miao, W., and Tchetgen Tchetgen, E. (2024). Semiparametric proximal causal inference. *Journal of the American Statistical Association*, 119(546):1348–1359.

- Delany-Moretlwe, S., Hughes, J. P., Bock, P., Ouma, S. G., Hunidzarira, P., Kalonji, D., Kayange, N., Makhema, J., Mandima, P., Mathew, C., et al. (2022). Cabotegravir for the prevention of hiv-1 in women: results from hptn 084, a phase 3, randomised clinical trial. *The Lancet*, 399(10337):1779–1789.
- Díaz, I., Williams, N., Hoffman, K. L., and Schenck, E. J. (2023). Nonparametric causal effects based on longitudinal modified treatment policies. *Journal of the American Statistical Association*, 118(542):846–857.
- Díaz Muñoz, I. and Van Der Laan, M. (2012). Population intervention causal effects based on stochastic interventions. *Biometrics*, 68(2):541–549.
- Dikkala, N., Lewis, G., Mackey, L., and Syrgkanis, V. (2020). Minimax estimation of conditional moment models. *Advances in Neural Information Processing Systems*, 33:12248–12262.
- Donnell, D., Gao, F., Hughes, J. P., Hanscom, B., Corey, L., Cohen, M. S., Edupuganti, S., Mgodi, N., Rees, H., Baeten, J. M., et al. (2023). Counterfactual estimation of efficacy against placebo for novel prep agents using external trial data: example of injectable cabotegravir and oral prep in women. *Journal of the International AIDS Society*, 26(6):e26118.
- Escanciano, J. C. and Li, W. (2013). On the identification of structural linear functionals. Technical report, cemmap working paper.
- FDA et al. (2016). Non-inferiority clinical trials to establish effectiveness guidance for industry. *US: Department of Health and Human Services*.
- Feng, S., Phillips, D. J., White, T., Sayal, H., Aley, P. K., Bibi, S., Dold, C., Fuskova, M.,

- Gilbert, S. C., Hirsch, I., et al. (2021). Correlates of protection against symptomatic and asymptomatic sars-cov-2 infection. *Nature medicine*, 27(11):2032–2040.
- Fleming, T. R. (2008). Current issues in non-inferiority trials. *Statistics in medicine*, 27(3):317–332.
- Fleming, T. R., Krause, P. R., Nason, M., Longini, I. M., and Henao-Restrepo, A.-M. M. (2021). Covid-19 vaccine trials: The use of active controls and non-inferiority studies. *Clinical Trials*, 18(3):335–342.
- Fleming, T. R., Odem-Davis, K., Rothmann, M. D., and Li Shen, Y. (2011). Some essential considerations in the design and conduct of non-inferiority trials. *Clinical Trials*, 8(4):432–439.
- Fleming, T. R. and Powers, J. H. (2012). Biomarkers and surrogate endpoints in clinical trials. *Statistics in medicine*, 31(25):2973–2984.
- Fong, Y., Huang, Y., Benkeser, D., Carpp, L. N., Áñez, G., Woo, W., McGarry, A., Dunkle, L. M., Cho, I., Houchens, C. R., et al. (2023). Immune correlates analysis of the prevent-19 covid-19 vaccine efficacy clinical trial. *Nature Communications*, 14(1):331.
- Fong, Y., McDermott, A. B., Benkeser, D., Roels, S., Stieh, D. J., Vandebosch, A., Le Gars, M., Van Roey, G. A., Houchens, C. R., Martins, K., et al. (2022). Immune correlates analysis of the ensemble single ad26. cov2. s dose vaccine efficacy clinical trial. *Nature Microbiology*, 7(12):1996–2010.
- Foster, D. J. and Syrgkanis, V. (2023). Orthogonal statistical learning. *The Annals of Statistics*, 51(3):879–908.

- Freyberger, J. and Horowitz, J. L. (2015). Identification and shape restrictions in nonparametric instrumental variables estimation. *Journal of Econometrics*, 189(1):41–53.
- Garreau, D., Jitkrittum, W., and Kanagawa, M. (2017). Large sample analysis of the median heuristic. *arXiv preprint arXiv:1707.07269*.
- Ghassami, A., Ying, A., Shpitser, I., and Tchetgen, E. T. (2022). Minimax kernel machine learning for a class of doubly robust functionals with application to proximal causal inference. In *International conference on artificial intelligence and statistics*, pages 7210–7239. PMLR.
- Gilbert, P. B., Donis, R. O., Koup, R. A., Fong, Y., Plotkin, S. A., and Follmann, D. (2022a). A covid-19 milestone attained—a correlate of protection for vaccines. *New England Journal of Medicine*, 387(24):2203–2206.
- Gilbert, P. B., Montefiori, D. C., McDermott, A. B., Fong, Y., Benkeser, D., Deng, W., Zhou, H., Houchens, C. R., Martins, K., Jayashankar, L., et al. (2022b). Immune correlates analysis of the mrna-1273 covid-19 vaccine efficacy clinical trial. *Science*, 375(6576):43–50.
- Gilead Sciences (2024). Gilead’s Twice-Yearly Lenacapavir for HIV Prevention Reduced HIV Infections by 96% and Demonstrated Superiority to Daily Truvada® in Second Pivotal Phase 3 Trial. Press release.
- Grant, R. M., Lama, J. R., Anderson, P. L., McMahan, V., Liu, A. Y., Vargas, L., Goicochea, P., Casapía, M., Guanira-Carranza, J. V., Ramirez-Cardich, M. E., et al. (2010). Preexposure chemoprophylaxis for hiv prevention in men who have sex with men. *New England Journal of Medicine*, 363(27):2587–2599.
- Haneuse, S. and Rotnitzky, A. (2013). Estimation of the effect of interventions that modify the received treatment. *Statistics in medicine*, 32(30):5260–5277.

- Hejazi, N. S., van der Laan, M. J., Janes, H. E., Gilbert, P. B., and Benkeser, D. C. (2021). Efficient nonparametric inference on the effects of stochastic interventions under two-phase sampling, with applications to vaccine efficacy trials. *Biometrics*, 77(4):1241–1253.
- Henderson, M., Schmidt, H.-M. A., Chitembo, L., Peralta, H., Alaama, A. S., Johnson, C., Baggaley, R., and Schaefer, R. (2023). The future of pre-exposure prophylaxis (prep) for hiv prevention: A global qualitative consultation on provider perspectives on new products and differentiated service delivery. *AIDS and Behavior*, pages 1–12.
- Hohage, T. (2002). Lecture notes on inverse problems. *University of Gottingen*, pages 14–21.
- Huang, Y., Hejazi, N. S., Blette, B., Carpp, L. N., Benkeser, D., Montefiori, D. C., McDermott, A. B., Fong, Y., Janes, H. E., Deng, W., et al. (2023). Stochastic interventional vaccine efficacy and principal surrogate analyses of antibody markers as correlates of protection against symptomatic covid-19 in the cove mrna-1273 trial. *Viruses*, 15(10):2029.
- Hung, H., Wang, S.-J., Tsong, Y., Lawrence, J., and O’Neil, R. T. (2003). Some fundamental issues with non-inferiority testing in active controlled trials. *Statistics in Medicine*, 22(2):213–225.
- Hung, H. J., Wang, S.-J., and O’Neill, R. (2007). Issues with statistical risks for testing methods in noninferiority trial without a placebo arm. *Journal of Biopharmaceutical Statistics*, 17(2):201–213.
- Kallus, N., Mao, X., and Uehara, M. (2021). Causal inference under unmeasured confounding with negative controls: A minimax learning approach. *arXiv preprint arXiv:2103.14029*.
- Kennedy, E. H. (2024). Semiparametric doubly robust targeted double machine learning: a review. *Handbook of statistical methods for precision medicine*, pages 207–236.

- Kress, R. (2014). *Linear Integral Equations*, volume 82 of *Applied Mathematical Sciences*. Springer, 3 edition.
- Landovitz, R. J., Donnell, D., Clement, M. E., Hanscom, B., Cottle, L., Coelho, L., Cabello, R., Chariyalertsak, S., Dunne, E. F., Frank, I., et al. (2021). Cabotegravir for hiv prevention in cisgender men and transgender women. *New England Journal of Medicine*, 385(7):595–608.
- Li, W., Miao, W., and Tchetgen Tchetgen, E. (2023). Non-parametric inference about mean functionals of non-ignorable non-response data without identifying the joint distribution. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(3):913–935.
- Liao, L., Chen, Y.-L., Yang, Z., Dai, B., Kolar, M., and Wang, Z. (2020). Provably efficient neural estimation of structural equation models: An adversarial approach. *Advances in Neural Information Processing Systems*, 33:8947–8958.
- Marrazzo, J. M., Ramjee, G., Richardson, B. A., Gomez, K., Mgodhi, N., Nair, G., Palanee, T., Nakabiito, C., Van Der Straten, A., Noguchi, L., et al. (2015). Tenofovir-based preexposure prophylaxis for hiv infection among african women. *New England Journal of Medicine*, 372(6):509–518.
- Mastouri, A., Zhu, Y., Gultchin, L., Korba, A., Silva, R., Kusner, M., Gretton, A., and Muandet, K. (2021). Proximal causal learning with kernels: Two-stage estimation and moment restriction. In *International conference on machine learning*, pages 7512–7523. PMLR.
- Miao, W., Geng, Z., and Tchetgen Tchetgen, E. J. (2018). Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*, 105(4):987–993.

- Miao, W., Liu, L., Li, Y., Tchetgen Tchetgen, E. J., and Geng, Z. (2024). Identification and semiparametric efficiency theory of nonignorable missing data with a shadow variable. *ACM/JMS Journal of Data Science*, 1(2):1–23.
- Miao, W. and Tchetgen Tchetgen, E. J. (2016). On varieties of doubly robust estimators under missingness not at random with a shadow variable. *Biometrika*, 103(2):475–482.
- Moyo, E., Murewanhema, G., Musuka, G., and Dzinamarira, T. (2022). Long-acting injectable drugs for hiv-1 pre-exposure prophylaxis: considerations for africa. *Tropical Medicine and Infectious Disease*, 7(8):154.
- Muandet, K., Jitkrittum, W., and Kübler, J. (2020a). Kernel conditional moment test via maximum moment restriction. In *Conference on Uncertainty in Artificial Intelligence*, pages 41–50. PMLR.
- Muandet, K., Mehrjou, A., Lee, S. K., and Raj, A. (2020b). Dual instrumental variable regression. *Advances in Neural Information Processing Systems*, 33:2710–2721.
- Murewanhema, G., Musuka, G., Moyo, P., Moyo, E., and Dzinamarira, T. (2022). Hiv and adolescent girls and young women in sub-saharan africa: A call for expedited action to reduce new infections. *IJID Regions*, 5:30–32.
- Newey, W. K. and Powell, J. L. (2003). Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578.
- Odem-Davis, K. and Fleming, T. R. (2013). Adjusting for unknown bias in noninferiority clinical trials. *Statistics in Biopharmaceutical Research*, 5(3):248–258.
- Park, C., Stensrud, M., and Tchetgen, E. T. (2024). Proximal causal inference for conditional separable effects. *arXiv preprint arXiv:2402.11020*.

- Paulsen, V. I. and Raghupathi, M. (2016). *An introduction to the theory of reproducing kernel Hilbert spaces*, volume 152. Cambridge university press.
- Plotkin, S. A. (2010). Correlates of protection induced by vaccination. *Clinical and vaccine immunology*, 17(7):1055–1065.
- Robins, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling*, 7(9-12):1393–1512.
- Robins, J., Li, L., Tchetgen, E., van der Vaart, A., et al. (2008). Higher order influence functions and minimax estimation of nonlinear functionals. In *Probability and statistics: essays in honor of David A. Freedman*, volume 2, pages 335–422. Institute of Mathematical Statistics.
- Rothmann, M., Li, N., Chen, G., Chi, G. Y., Temple, R., and Tsou, H.-H. (2003). Design and analysis of non-inferiority mortality trials in oncology. *Statistics in medicine*, 22(2):239–264.
- Rothmann, M. D., Wiens, B. L., and Chan, I. S. (2011). *Design and analysis of non-inferiority trials*. CRC press.
- Sadoff, J., Gray, G., Vandebosch, A., Cárdenas, V., Shukarev, G., Grinsztejn, B., Goepfert, P. A., Truyers, C., Van Dromme, I., Spiessens, B., et al. (2022). Final analysis of efficacy and safety of single-dose ad26. cov2. s. *New England Journal of Medicine*, 386(9):847–860.
- Santos, A. (2011). Instrumental variable methods for recovering continuous linear functionals. *Journal of Econometrics*, 161(2):129–146.

- Schölkopf, B., Herbrich, R., and Smola, A. J. (2001). A generalized representer theorem. In *International conference on computational learning theory*, pages 416–426. Springer.
- Severini, T. A. and Tripathi, G. (2012). Efficiency bounds for estimating linear functionals of nonparametric regression models with endogenous regressors. *Journal of Econometrics*, 170(2):491–498.
- Shi, X., Miao, W., Nelson, J. C., and Tchetgen Tchetgen, E. J. (2020). Multiply robust causal inference with double-negative control adjustment for categorical unmeasured confounding. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(2):521–540.
- Sholukh, A. M., Fiore-Gartland, A., Ford, E. S., Hou, Y. J., Tse, V., Kaiser, H., Park, A., Lempp, F. A., Germain, R. S., Bossard, E., et al. (2020). Evaluation of sars-cov-2 neutralization assays for antibody monitoring in natural infection and vaccine trials. *medRxiv*, pages 2020–12.
- Snapinn, S. and Jiang, Q. (2008). Controlling the type 1 error rate in non-inferiority trials. *Statistics in Medicine*, 27(3):371–381.
- Steinwart, I. and Christmann, A. (2008). *Support vector machines*. Springer.
- Tchetgen, E. J. T., Ying, A., Cui, Y., Shi, X., and Miao, W. (2020). An introduction to proximal causal learning. *arXiv preprint arXiv:2009.10982*.
- Thigpen, M. C., Kebaabetswe, P. M., Paxton, L. A., Smith, D. K., Rose, C. E., Segolodi, T. M., Henderson, F. L., Pathak, S. R., Soud, F. A., Chillag, K. L., et al. (2012). Antiretroviral preexposure prophylaxis for heterosexual hiv transmission in botswana. *New England Journal of Medicine*, 367(5):423–434.
- UNAIDS (2023). The path that ends aids: 2023 unaids global aids update.

- Van Damme, L., Corneli, A., Ahmed, K., Agot, K., Lombaard, J., Kapiga, S., Malahleha, M., Owino, F., Manongi, R., Onyango, J., et al. (2012). Preexposure prophylaxis for hiv infection among african women. *New England Journal of Medicine*, 367(5):411–422.
- van der Straten, A., Stadler, J., Luecke, E., Laborde, N., Hartmann, M., Montgomery, E. T., and Team, V.-C. S. (2014). Perspectives on use of oral and vaginal antiretrovirals for hiv prevention: the voice-c qualitative study in johannesburg, south africa. *Journal of the International AIDS Society*, 17:19146.
- Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press.
- Walker, B. D. (2021). The amp trials—a glass half full. *N Engl J Med*, 384(11):1068–9.
- Wang, S.-J., Hung, H. J., and Tsong, Y. (2002). Utility and pitfalls of some statistical methods in active controlled clinical trials. *Controlled Clinical Trials*, 23(1):15–28.
- Zhang, R., Imaizumi, M., Schölkopf, B., and Muandet, K. (2023). Instrumental variable regression via kernel maximum moment loss. *Journal of Causal Inference*, 11(1):20220073.