



HAL
open science

Targeted learning in Big Data : bridging data-adaptive estimation and statistical inference

Wenjing Zheng

► **To cite this version:**

Wenjing Zheng. Targeted learning in Big Data : bridging data-adaptive estimation and statistical inference. General Mathematics [math.GM]. Université Sorbonne Paris Cité, 2016. English. NNT : 2016USPCB044 . tel-01730786

HAL Id: tel-01730786

<https://theses.hal.science/tel-01730786>

Submitted on 13 Mar 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



UNIVERSITÉ PARIS DESCARTES

Laboratoire MAP5 UMR CNRS 8145

École doctorale 386 : Sciences Mathématiques de Paris Centre

THÈSE

Pour obtenir le grade de

DOCTEUR de l'UNIVERSITÉ PARIS DESCARTES

Spécialité : Mathématiques appliquées

Présentée par

Wenjing ZHENG

**Apprentissage ciblé et Big Data: contribution à la
réconciliation de l'estimation adaptative et de l'inférence
statistique**

soutenue le 21 juillet 2016 devant le jury composé de

Antoine Chambaz	Université Paris Nanterre	Co-Directeur de thèse
Mark van der Laan	University of California, Berkeley, USA	Co-Directeur de thèse
Catherine Matias	Université Pierre et Marie Curie	Examinatrice
Cristina Butucea	Université Paris-Est Marne-la-Vallée	Examinatrice
Rodolphe Thiébaud	Université de Bordeaux	Examineur
Erica Moodie	McGill University, Canada	Rapportrice
Stijn Vansteelandt	Universiteit Gent, Belgium	Rapporteur

Apprentissage ciblé et Big Data: contribution à la réconciliation de l'estimation adaptative et de l'inférence statistique

Résumé

Cette thèse porte sur le développement de méthodes semi-paramétriques robustes pour l'inférence de paramètres complexes émergeant à l'interface de l'inférence causale et la biostatistique. Ses motivations sont les applications à la recherche épidémiologique et médicale à l'ère des Big Data. Nous abordons plus particulièrement deux défis statistiques pour réconcilier, dans chaque contexte, estimation adaptative et inférence statistique.

Le premier défi concerne la maximisation de l'information tirée d'essais contrôlés randomisés (ECRs) grâce à la conception d'essais adaptatifs. Nous présentons un cadre théorique pour la construction et l'analyse d'ECRs groupes-séquentiels, réponses-adaptatifs et ajustés aux covariable (traduction de l'expression anglaise « group-sequential, response-adaptive, covariate-adjusted », d'où l'acronyme CARA) qui permettent le recours à des procédures adaptatives d'estimation à la fois pour la construction dynamique des schémas de randomisation et pour l'estimation du modèle de réponse conditionnelle. Ce cadre enrichit la littérature existante sur les ECRs CARA notamment parce que l'estimation des effets est garantie robuste même lorsque les modèles sur lesquels s'appuient les procédures adaptatives d'estimation sont mal spécifiés.

Le second défi concerne la mise au point et l'étude asymptotique d'une procédure inférentielle semi-paramétrique avec estimation adaptative des paramètres de nuisance. A titre d'exemple, nous choisissons comme paramètre d'intérêt la différence des risques marginaux pour un traitement binaire. Nous proposons une version cross-validée du principe d'inférence par minimisation ciblée de pertes (« Cross-validated Targeted Minimum Loss Estimation » en anglais, d'où l'acronyme CV-TMLE) qui, comme son nom le suggère, marie la procédure TMLE classique et le principe de la validation croisée. L'estimateur CV-TMLE ainsi élaboré hérite de la propriété typique de double-robustesse et aussi des propriétés d'efficacité du TMLE classique. De façon remarquable, le CV-TMLE est linéairement asymptotique sous des conditions minimales, sans recourir aux conditions de type Donsker.

Remerciements

La réalisation de cette thèse de doctorat a été rendue possible grâce à la générosité, à l'aide et au soutien de nombreuses personnes et organisations. Je souhaite tout d'abord remercier le MAP5 (UMR CNR 8145, Université Paris Descartes) pour son accueil lors de ma longue visite scientifique parrainée par l'Ambassade des Etats-Unis à Washington, à laquelle j'exprime également ma gratitude, dans le cadre d'une Bourse STEM Chateaubriand. Je veux aussi remercier le projet SPADRO (ANR-13-BS01-0005) de l'Agence Nationale de la Recherche, animé par Antoine Chambaz et Aurélien Garivier, qui a financé plusieurs de mes missions en France.

Je tiens à exprimer ma profonde reconnaissance aux examinateurs de ma thèse, les professeurs Erica Moodie et Stijn Vansteelandt. Leurs commentaires éclairés, approfondis, constructifs et bienveillants m'ont poussée à développer davantage mes réflexions sur les sujets de ma thèse et ont conduit à de nombreuses améliorations du manuscrit. Je souhaite également remercier les professeurs Catherine Matias, Cristina Butucea et Rodolphe Thiébaud de me faire l'honneur de constituer mon jury de soutenance.

J'adresse ma plus profonde gratitude à mes co-directeurs de thèse, Mark van der Laan et Antoine Chambaz, et à ma mentor de carrière, Maya Petersen. Les sentiments les plus sincères sont souvent les plus difficiles à exprimer, aussi je souhaite que ma reconnaissance ne soit pas mesurée à l'aune seule des mots que je réussirai à aligner. Je serai éternellement reconnaissante à Mark van der Laan d'avoir été un directeur si passionné, patient, excitant, stimulant et inspirant, m'aidant à toujours repousser mes limites. Ce fut un privilège d'apprendre les statistiques de lui, et une toute nouvelle façon de penser ; j'en récolterai les fruits pour le reste de ma vie professionnelle et personnelle. Son éthique de travail exemplaire, sa passion, sa résilience seront une inspiration pour accéder au même standard. Mes plus sincères remerciements vont également à Antoine Chambaz, pour son dévouement, sa passion, sa patience et ses encouragements et soutien inlassables, et pour être l'un des collaborateurs et mentors les plus amusants, passionnants et inspirants avec qui j'aie eu l'honneur de travailler. Sa stimulation intellectuelle continue et son attitude m'ont aidée à aller au-delà de ce que je

croyais être mes limites. Je suis également reconnaissante à Antoine pour la précieuse occasion de travailler avec lui à Paris, qui est l'un des plus beaux souvenirs de ma carrière d'études supérieures. Je voudrais aussi exprimer ma profonde gratitude à Julie, Lou, Fausto et Claire, pour m'avoir accueillie chez eux pendant mon séjour en France; l'hiver parisien épique ne pouvait pas rivaliser avec la chaleur de leur belle famille. Je suis également reconnaissants envers le professeur Maya Petersen de UC Berkeley pour son soutien permanent et ses conseils, pour son enthousiasme contagieux, et pour l'aide qu'elle m'a si généreusement prodiguée alors que je commençais ma carrière une fois obtenu mon diplôme. En apprenant de Maya l'inférence causale et la résolution de problèmes appliqués, j'ai gagné une toute nouvelle perspective et j'ai compris comment, en statisticienne ou statisticien, nous pouvons contribuer directement et immédiatement à la résolution de problèmes issus du monde réel, et ce que sont les responsabilités inhérentes.

Enfin, le long voyage de cette formation doctorale a été alimenté par le soutien indéfectible et inébranlable de ma famille. Je suis très reconnaissante à mon merveilleux mari Karén et à notre beau fils Ruben pour leur compréhension et pour leur aide constante à bien garder en perspective ce que sont les choses les plus importantes dans la vie.

Targeted Learning in Big Data: Bridging data-adaptive estimation and statistical inference

Summary

This dissertation focuses on developing robust semiparametric methods for complex parameters that emerge at the interface of causal inference and biostatistics, with applications to epidemiological and medical research in the era of Big Data. Specifically, we address two statistical challenges that arise in bridging the disconnect between data-adaptive estimation and statistical inference. The first challenge arises in maximizing information learned from Randomized Control Trials (RCT) through the use of adaptive trial designs. We present a framework to construct and analyze group sequential covariate-adjusted response-adaptive (CARA) RCTs that admits the use of data-adaptive approaches in constructing the randomization schemes and in estimating the conditional response model. This framework adds to the existing literature on CARA RCTs by allowing flexible options in both their design and analysis and by providing robust effect estimates even under model mis-specifications. The second challenge arises from obtaining a Central Limit Theorem when data-adaptive estimation is used to estimate the nuisance parameters. We consider as target parameter of interest the marginal risk difference of the outcome under a binary treatment, and propose a Cross-validated Targeted Minimum Loss Estimator (TMLE), which augments the classical TMLE with a sample-splitting procedure. The proposed Cross-Validated TMLE (CV-TMLE) inherits the double robustness properties and efficiency properties of the classical TMLE, and achieves asymptotic linearity at minimal conditions by avoiding the Donsker class condition.

Acknowledgments

The completion of this doctoral work was only possible thanks to the generosity, help and support of many organizations and individuals. I must thank Laboratoire MAP5 for hosting my fellowship and for granting me this opportunity. I am also thankful to the Chateaubriand STEM fellowship, whose sponsorship for my adaptive design project has made a large part of this thesis possible, and to the support of ANR-SPADRO, whose grant to Professor Antoine Chambaz has made my travels to France possible. I am deeply grateful to the reviewers Professors Erica Moodie and Stijn Vansteelandt for their time, and for their thorough, insightful, constructive and thoughtful comments that have pushed me to further develop my thoughts on the thesis topics and have led to many improvements in this thesis. I am equally thankful to the defense committee members Professors Catherine Matias, Cristina Butucea and Rodolphe Thiébaud, for their generosity with their time and support that make this defense possible.

I owe my deepest gratitude to my thesis mentors Mark van der Laan and Antoine Chambaz, and my career mentor Maya Petersen. But the most heartfelt sentiments are often the most inarticulate ones; hard as I try to express these here, my gratitude to them cannot be gauged by the limitations of my words. I am eternally grateful to Mark van der Laan for being a passionate, patient, exciting, challenging and inspiring mentor, who was always pushing the limits of my abilities. It was a privilege to learn statistics and a whole new way of thinking from him, benefits of which I will reap for the rest of my professional and personal life. In addition to fostering my intellectual growth, his immaculate work ethics, drive, and resilience constantly inspire me to hold myself to the same standards. My deepest thanks also goes to Antoine Chambaz, for his dedication, passion, patience and tireless encouragement and support, and for being one of the most fun, exciting and inspiring collaborators and mentors I have had the honor to work with. His continuous intellectual stimulation and nurturing attitude constantly helped me go beyond the boundaries of my own expectations. I am also thankful to Antoine for the precious opportunity to work with him in Paris, which was one of the most wonderful memories of my graduate career. In addition to Antoine, I would like to extend my deep gratitude to

Julie, Lou, Fausto and Claire, for welcoming me into their home during my stay in France; that epic Parisian winter was no match for the warmth of their beautiful family. I am equally grateful to Professor Maya Petersen of UC Berkeley for her continuing support and guidance, for her contagious enthusiasm, and for her generous guidance in helping me navigate a career path post graduation. Through learning causal inference and solving application problems from Maya, I gained a whole new perspective and saw how, as a statistician, one can have direct immediate contribution to, and inherent responsibility in, solving real world problems.

Finally, the long trek of this doctoral training was fueled by the untiring and steadfast support of my family. I am very grateful to my wonderful husband, Karén and our beautiful son Ruben for their understanding and for helping me keep in perspective the most important things in life.

Contents

Contents

1	Introduction	1
1.1	Background and Motivation	1
2	When Adaptive CARA RCT meets Data-Adaptive Estimation: Targeted Maximum Likelihood Estimation for Adaptive Randomized Controlled Trial Designs	4
2.1	Introduction	4
2.2	Targeted inference based on data adaptively drawn from a CARA RCT using loss-based estimation	9
2.3	Comments on Section 2.2	15
2.4	Building blocks of Theorem 2.1	20
2.5	Example: targeted LASSO-based CARA RCT	22
2.6	Simulation study	24
2.7	Discussion	27
2.8	Appendix	29
3	Asymptotic Theory for Cross-validated Targeted Maximum Likelihood Estimation	54
3.1	Introduction	54
3.2	The TMLE using V-fold sample splitting for initial estimator.	55
3.3	Asymptotics for the one-step cross-validated TMLE	59
3.4	Application of Theorem 3.2 to estimation of additive causal effect in non-parametric model	71
3.5	The iterative targeted MLE using V-fold sample splitting.	89
3.6	Concluding remarks.	94
3.7	Appendix	94

4	Bibliography	108
4.1	Bibliography	108

Chapter 1

Introduction

1.1 Background and Motivation

The current Big Data revolution is marked by the unprecedented volume, diversity and velocity of data being collected across public and private sectors. This, coupled with dramatic increases in computing power, presents revolutionary opportunities to improve population and individual-level well-being. This deluge of data can yield real insight into today's most critical social and health problems across the globe, but statistical advances from the traditional paradigms are needed to fully realize its potential. Traditional statistical approaches that rely on finite-dimensional (parametric) modeling of the data-generating process cannot keep up with the increasing complexity and diversity of data, missing real opportunities to uncover previously unknown interactions and relations. While recent advances in machine learning methods inject the much needed flexibility to maximize the predictive power of Big Data, these methods often lack inferential capability as they do not provide formal statistical inferences on the underlying data generating distribution, and hence fail to produce actionable information to inform treatment development or policy planning. This disconnect between breakthrough in machine learning methods and the inferential power of classical statistics translates into missed potentials to deliver solutions for today's most urgent social and health problems. In this dissertation, we address two statistical challenges that arise in bridging this disconnect between data-adaptive prediction and statistical inference.

The first challenge arises in maximizing information learned from Randomized Control Trials (RCT). RCTs are the gold standard for causal effect inference and are ubiquitous in health and social science research. Traditionally, the designs for these trials are fixed, in the sense that all key aspects of the trial are set before the start of the data collection, usually based on assumptions that are yet unsure at the design stage. Adaptive trial de-

sign methods, by contrast, allows pre-specified modifications of the ongoing trial based on accruing data, while preserving the validity and integrity of the trial. This flexibility potentially translates into more tailored studies: the study could be more efficient, *e.g.*, have shorter duration, or involve fewer subjects; the study could have greater chance to answer the clinical questions of interest, *e.g.*, detect a treatment effect if one exists, or gather broader dose-response information. For this reason, adaptive trial designs have garnered growing attention in recent years. We focus here on the study of the so-called covariate-adjusted response-adaptive (CARA) randomized controlled trials (RCTs). In a CARA RCT, the randomization schemes are allowed to depend on the patient's pre-treatment covariates, and the investigators can adjust/adapt the randomization schemes during the course of the trial based on accruing information, including previous responses, in order to meet the pre-specified optimal design criteria. Inherently, since the randomization scheme is response-adaptive, how successfully it approximates the desired design optimality would critically depend on how well we can model the response of interest. In a heterogeneous patient population, responses may be affected by a wide array of baseline individual characteristics in manners unknown to the investigators. Therefore, for the goal of achieving design optimality without compromising trial integrity, we need to 1) do the best we can in estimating the outcome model, 2) produce unbiased estimates of the primary study parameter, regardless of the specification of outcome model, thus preserving the robust hallmark of an RCT. This requires a design and analysis framework to that allows robust inference of the study parameter of interest with data-adaptive estimators. The statistical challenge here, compare to classical fixed design, lies in the non-independently and non-identically distributed nature of the observed data. As a result, many empirical process theories essential to robust estimation and inferences under independent and identically distributed (i.i.d.) observations need to be extended to the CARA RCT sampling scheme. Chapter 2 proposes a framework for design and analysis of CARA RCT that incorporates the use of data-adaptive estimators (*e.g.* machine learning techniques) in modeling the response and in estimating the optimal design, and, using the Targeted Minimum Loss Estimation methodology (TMLE, van der Laan and Rubin 2006), provides robust inference of the study parameter under an arbitrarily misspecified response model. We call this framework a targeted CARA RCT in the sense that it targets both the unknown optimal randomization scheme by data-adaptively modifying the randomization schemes using accruing data, and the study parameter of interest by providing unbiased effect estimate.

The second challenge this dissertation addresses arises from obtaining a Central Limit Theorem when data-adaptive estimation is used to estimate the nuisance parameters. Consider a target parameter of interest that is a path-wise differentiable function of the data-generating distribution. Given an asymptotically linear estimator of this target parameter, we can use Central Limit Theorem to obtain an Influence Curve-based variance estimate

for this estimator, and therefore produce valid confidence intervals for formal hypothesis testing. In many applications, the asymptotic linearity relies on Donsker class conditions on the empirical processes related to the estimators of the nuisance parameters. These conditions effectively restrict how data-adaptive such nuisance parameter estimates can be, and thereby excluding the use of a large class of machine learning techniques that may help reduce the bias of the target parameter estimate. In Chapter 3, we consider as target parameter of interest the marginal risk difference of the outcome under a binary treatment, and propose a Cross-validated Targeted Minimum Loss Estimator, which augments the classical TMLE with a sample-splitting procedure. The classical TMLE provides a doubly robust substitution estimator for the target parameter of interest by updating the initial estimators of the response model using a least favorable submodel. TMLE produces unbiased estimate of the target parameter if the response or the treatment assignment are consistently estimated. It is asymptotically linear if Donsker class conditions are satisfied, which case if all nuisance parameters are correct, it is also asymptotically efficient. The proposed Cross-Validated TMLE (CV-TMLE) inherits the double robustness properties and efficiency properties, and achieves asymptotic linearity at minimal conditions by avoiding the Donsker class condition.

Chapter 2

When Adaptive CARA RCT meets Data-Adaptive Estimation: Targeted Maximum Likelihood Estimation for Adaptive Randomized Controlled Trial Designs

2.1 Introduction

Covariate-adjusted, response-adaptive randomized clinical trials

Adaptive clinical trial design methods have garnered growing attention in recent years. In a fixed trial design, all key aspects of the trials are set before the start of the data collection, usually based on assumptions that are yet unsure at the design stage. By contrast, an adaptive trial design allows pre-specified modifications of the ongoing trial based on accruing data, while preserving the validity and integrity of the trial. This flexibility potentially translates into more tailored studies. The study could be more efficient, *e.g.*, have shorter duration, or involve fewer subjects. The study could have greater chance to answer the clinical questions of interest, *e.g.*, detect a treatment effect if one exists, or gather broader dose-response information.

Once they have defined the *primary study objective* of the trial (*e.g.*, testing the effect of a treatment), the investigators may wish to accommodate additional *design objectives* (*e.g.*, minimizing sample size or exposure of patients to inferior treatment) without compromising the trial. To do this, they may use an adaptive randomization trial design. We

focus here on the study of the so-called covariate-adjusted response-adaptive (CARA) randomized controlled trials (RCTs). In a CARA RCT, the randomization schemes are allowed to depend on the patient's pre-treatment covariates, and the investigators can adjust the randomization schemes during the course of the trial based on accruing information, including previous responses, in order to meet the pre-specified design objectives. Such adjustments take place at interim time points given by sequential inclusion of blocks of c patients, where $c \geq 1$ is a pre-specified integer. We consider the case of $c = 1$ for simplicity of exposition, though the discussions generalize to any $c > 1$.

The trial protocol pre-specifies the observed data structure, scientific parameters of interest (the primary study objective), analysis methods, and a criterion characterizing an optimal randomization scheme (the design objective). Here, baseline covariates and a primary response of interest are measured on each patient. The primary study objective is the marginal treatment effect. The design objective is captured by an optimality criterion that is a function of the unknown conditional response.

Contrary to a fixed design RCT, a CARA RCT produces non-independent and non-identically distributed observations, therein lie the subtleties in its theoretical study. Traditionally, covariate-adjusted analysis of a fixed design RCT is carried out using a parametric model for the conditional response (or distribution) given treatment and covariates. Under correct specification, the maximum likelihood estimator for this model is consistent and asymptotically Gaussian. The extension of this *non-robust inference* to CARA RCTs has been established and discussed in (Zhang, Hu, Cheung, and Chan, 2007) and (Hu and Rosenberger, 2006). A recent development in the analysis of fixed design RCTs is the use of doubly robust methods like the targeted minimum loss estimation (TMLE, van der Laan and Rubin 2006) to obtain consistent, asymptotically Gaussian estimators under arbitrarily mis-specified models (Moore and van der Laan, 2009, Rosenblum, 2011). A first extension of this *robust inference* to CARA RCTs has been proposed by Chambaz and van der Laan (2013). They showed that, when the treatment assignment is conditioned only on a discrete summary measure of the covariates, it is possible to derive a consistent and asymptotically Gaussian estimator of the study parameter which is robust to mis-specification of an arbitrary parametric response model.

Despite the above developments, several gaps remain to be addressed to fully realize the promise of CARA RCTs. We focus on two of them. Firstly, because the robust inference provided by Chambaz and van der Laan (2013) relies on assigning treatment based on discrete covariate summaries, its application is perhaps limited in real-life RCTs where response to treatment may be correlated with a large number of a patient's baseline characteristics, some of which being continuous. Secondly, even though under robust inference, the choice of the response model does not compromise consistent estimation of the study parameter, it may still affect the estimation of the optimal randomization scheme. Specifically, since the randomization scheme is response-adaptive, a more data-adaptive

estimator of the conditional response model can more effectively steer the randomization schemes towards the unknown optimal randomization scheme. Moreover, since a patient's primary response is often correlated with many individual characteristics, greater latitude in adjusting for these baseline covariates, in both treatment assignment and conditional response estimation, allows the investigators to better account for heterogeneity in the patients population. Traditional parametric regression techniques are often too restrictive in such a high-dimensional scenario. While the use of data-adaptive techniques is very common in the independent and identically distributed (i.i.d.) context, its applicability in an adaptive RCT remains rather uncharted.

In this article, we aim to bridge the two aforementioned gaps in the study of CARA RCTs. Firstly, we achieve robust inference of the study parameter without restrictions on the covariate measures used in the treatment randomization. Secondly, we adopt the use of loss-based data-adaptive estimation over general classes of functions (which may change with sample size) in constructing the treatment randomization schemes and in predicting the unknown conditional response. This allows one to target general randomization optimality criteria that may not have a closed form solution, and it may potentially improve the estimation of the unknown optimal randomization schemes. We establish that, under appropriate entropy conditions on the classes of functions, the resulting sequence of randomization schemes converges to a fixed randomization scheme, and the proposed estimator is consistent (even under a mis-specified response model), asymptotically Gaussian, and gives rise to valid confidence intervals of given asymptotic levels. Moreover, the limiting randomization scheme coincides with the unknown optimal randomization scheme when, simultaneously, the response model is correctly specified and the optimal randomization scheme belongs to the limit of the user-supplied classes of randomization schemes. Our theoretical results benefit from recent advances in maximal inequalities for martingales by van Handel (2011).

For concreteness, our parameter of interest here is the marginal risk difference, ψ_0 , and our design objective is to maximize the efficiency of the study (*i.e.*, to reach a valid result using as few blocks of patients as possible). As we shall see, the optimal randomization scheme is, in this case, the so-called covariate-adjusted Neyman design, which minimizes the Cramér-Rao lower bound on the asymptotic variances of a large class of estimators of ψ_0 . We emphasize that the results presented here are not limited to the marginal risk difference or the Neyman design, and can be easily modified to other study objectives/effect measures and other design objectives/optimality criteria.

To illustrate the proposed framework, we consider the LASSO to estimate the conditional response given treatment and baseline covariates and to target the unknown optimal randomization scheme. This example essentially encompasses the parametric approach in (Chambaz and van der Laan, 2013) as a special case. The asymptotic results ensue under minimal conditions on the smoothness of the LASSO basis functions. The performance of

the procedure is evaluated in a simulation study.

In the next section, we give a bird's eye view of the literature on CARA RCTs and put our contribution in context.

Literature review

Adaptive randomization has a long history that can be traced back to the 1930s. We refer to (Rosenberger, 1996, Rosenberger, Sverdlov, and Hu, 2012), (Hu and Rosenberger, 2006, Section 1.2) and (Jennison and Turnbull, 2000, Section 17.4) for a comprehensive historical perspective. Many articles are devoted to the study of response-adaptive randomizations, which select current treatment probabilities based on responses of previous patients (but not on the covariates of the current patients). We summarize some representative works below, but refer to (Hu and Rosenberger, 2006, Chambaz and van der Laan, 2011b, Rosenberger et al., 2012) for a bibliography on that topic. The first methods are based on urn-models (e.g. Wei and Durham (1978), Ivanova (2003)). There, treatment allocation is represented by drawing balls of different colors from an urn, and the urn composition is updated based on accruing responses, with the ethical goal of assigning most patients to the superior treatment arm. Since there is no formal criterion governing how skewed the treatment allocation should be, significant loss of power can arise when the effect size between treatment arms is large (Rosenberger and Hu, 2004). A formal "optimal allocation approach" was proposed by Hu and Rosenberger (2003). There, an optimal allocation is defined as a solution to a possibly constrained optimization problem, such as minimizing sample size (yielding the so-called Neyman allocation), or minimizing failure while preserving power. This optimal allocation is a function of unknown parameters of the conditional response, which are estimated using a parametric model based on available responses. Consistency and asymptotic normality of the maximum likelihood estimator for this model were established in Hu and Rosenberger (2006).

In a heterogeneous population where response is often correlated with the patient's individual characteristics, covariates are often accounted for in treatment assignment. CARA randomization extends response-adaptive randomization to tackle heterogeneity by dynamically calculating the allocation probabilities based on previous responses and current and past values of certain covariates. Compared to the broader literature on response-adaptive randomization, the advances in CARA randomization are relatively recent, but growing steadily. Among the first approaches, (Rosenberger, Vidyashankar, and Agarwal, 2001, Bandyopadhyay and Biswas, 2001) considered allocations that are proportional to the covariate-adjusted treatment difference, which is estimated using generalized linear models for the conditional response. Though these procedures are not defined based on formal optimality criteria, their general goal is to allocate more patients to their corresponding superior treatment arm. Atkinson and Biswas (2005) presented a biased-coin

design with skewed allocation, which is determined by sequentially maximizing a function that combines the variance of the parameter estimator, based on a Gaussian linear model for the conditional response, and the conditional treatment effect given covariates. Up till here, very little work had been devoted to asymptotic properties of CARA designs. Subsequently, Zhang et al. (2007) established the asymptotic theory for CARA designs converging to a given target covariate-adjusted allocation function when the conditional responses follow a parametric model. Zhang and Hu (2009) proposed a covariate-adjusted doubly-adaptive biased coin design whose asymptotic variance achieves the efficiency bound. In these optimal allocation approaches, the challenge remains that the explicit form of the target covariate-adjusted allocation function is not known. To overcome this, it has often been derived as a covariate-adjusted version of the optimal allocation from a framework with no covariates (Rosenberger et al., 2012). Chang and Park (2013) proposed a sequential estimation of CARA designs under generalized linear models for the conditional response. This procedure allocates treatment based on the patients' baseline covariates, accruing information and sequential estimates of the treatment effect. It uses a stopping rule that depends on the observed Fisher information. With regard to hypotheses testing, Shao, Yu, and Zhong (2010), Shao and Yu (2013) provided asymptotic results for valid tests under generalized linear models for the conditional responses in the context of covariate-adaptive randomization. Most recently, progress has also been made in CARA designs in the longitudinal settings, see for example (Biswas, Bhattacharya, and Park, 2014, Huang, Liu, and Hu, 2013, Sverdlov, Rosenberger, and Ryznik, 2013).

The above contributions have established the validity of statistical inference for CARA RCTs under a correctly specified model, thus extending many of the classical non-robust inference methods from the fixed design setting into the CARA setting. Doubly robust approaches like TMLE allow to go beyond correctly specified models by leveraging the known treatment randomization to provide the necessary bias reduction over the misspecified response model. Moore and van der Laan (2009), Rosenblum (2011) address the fixed design setting and Chambaz and van der Laan (2013) provide the first extension to the adaptive design setting.

Finiteness conditions were at the core of (Zhang et al., 2007) (correctly specified parametric response model) and (Chambaz and van der Laan, 2013) (arbitrary parametric response model and treatment assignment based on discrete covariates). They were instrumental in the asymptotic study based on Taylor approximations. Recent advances by van Handel (2011) on maximal deviation bounds for martingales allow us to apply more general empirical processes techniques, thus opening the door for the use of data-adaptive estimators to target the optimal randomization scheme while preserving valid inference. More specifically, we extend the robust inference framework of Chambaz and van der Laan (2013) to allow for the use of general classes of conditional response estimators and randomization schemes. Moreover, we adopt a loss-based approach to defining and targeting

the optimal randomization scheme, thereby also extending applicability of CARA RCT to optimal randomization criteria that may not have a closed form solution.

Organization of the article

The remainder of this article is organized as follows. Section 2.2 presents the statistical challenges, and describes the proposed targeted, adaptive sampling scheme and inference method to address them. The section also states our main assumptions and principal result. Section 2.3 provides contextual comments for content of Section 2.2. Section 2.4 presents the building blocks of the main result, therefore shedding light on the inner mechanism of the procedure. Section 2.5 illustrates the procedure using the LASSO methodology both to target the optimal randomization scheme and to estimate the conditional response given baseline covariates and treatment. The performance of the LASSO-based CARA RCT is assessed in a simulation study in Section 2.6. The article closes on a discussion in Section 2.7. All proofs and some useful, technical results are gathered in the appendix.

2.2 Targeted inference based on data adaptively drawn from a CARA RCT using loss-based estimation

At sample size n , we will have observed the ordered vector $\mathbf{O}_n \equiv (O_1, \dots, O_n)$, with convention $O_0 \equiv \emptyset$. For every $1 \leq i \leq n$, the data structure O_i writes as $O_i \equiv (W_i, A_i, Y_i)$. Here, $W_i \in \mathcal{W}$ consists of the baseline covariates (some of which may be continuous) of the i th patient, $A_i \in \mathcal{A} \equiv \{0, 1\}$ is the binary treatment of interest assigned to her, and $Y_i \in \mathcal{Y}$ is her primary response of interest. We assume that the outcome space $\mathcal{O} \equiv \mathcal{W} \times \mathcal{A} \times \mathcal{Y}$ is bounded. Without loss of generality, we may then assume that $\mathcal{Y} \equiv (0, 1)$, *i.e.*, that the responses are between and bounded away from 0 and 1.

Section 2.2 presents the target statistical parameter and optimal randomization scheme. It also lays out the foundations to describe the proposed CARA RCT. The description is completed in Sections 2.2 and 2.2, where we present our adaptive sampling scheme and targeted minimum loss estimator. Section 2.2 states our main assumptions and result.

Likelihood, model, statistical parameter, optimal randomization scheme

Let μ_W be a measure on \mathcal{W} equipped with a σ -field, $\mu_A = \text{Dirac}(0) + \text{Dirac}(1)$ be a measure on \mathcal{A} equipped with its σ -field, and μ_Y be the Lebesgue measure on \mathcal{Y} equipped with the Borel σ -field. Define $\mu \equiv \mu_W \otimes \mu_A \otimes \mu_Y$, a measure on \mathcal{O} equipped with the product

of the above σ -fields. In an RCT, the unknown, true likelihood of \mathbf{O}_n with respect to (wrt) $\mu^{\otimes n}$ is given by the following factorization of the density of \mathbf{O}_n wrt $\mu^{\otimes n}$:

$$\begin{aligned} \mathcal{L}_{\mathbf{f}_0, \mathbf{g}_n}(\mathbf{O}_n) &\equiv \prod_{i=1}^n \mathbf{f}_{W,0}(W_i) \times (A_i g_i(1|W_i) + (1 - A_i) g_i(0|W_i)) \times \mathbf{f}_{Y,0}(Y_i|A_i, W_i) \\ &= \prod_{i=1}^n \mathbf{f}_{W,0}(W_i) \times g_i(A_i|W_i) \times \mathbf{f}_{Y,0}(Y_i|A_i, W_i), \end{aligned} \quad (2.1)$$

where (i) $w \mapsto \mathbf{f}_{W,0}(w)$ is the density wrt μ_W of a true, unknown law $Q_{W,0}$ on \mathscr{W} (that we assume being dominated by μ_W), (ii) $\{y \mapsto \mathbf{f}_{Y,0}(y|a, w) : (a, w) \in \mathscr{A} \times \mathscr{W}\}$ is the collection of the conditional densities $y \mapsto \mathbf{f}_{Y,0}(y|a, w)$ wrt μ_Y of true, unknown laws on \mathscr{Y} indexed by (a, w) (that we assume being all dominated by μ_Y), (iii) $g_i(1|W_i)$ is the known (given by user) conditional probability that $A_i = 1$ given W_i , and (iv) $\mathbf{g}_n \equiv (g_1, \dots, g_n)$, the ordered vector of the n first randomization schemes. One reads in (2.1) (i) that W_1, \dots, W_n are independently sampled from $Q_{W,0}$, (ii) that Y_1, \dots, Y_n are conditionally sampled from $\mathbf{f}_{Y,0}(\cdot|A_1, W_1)\mu_Y, \dots, \mathbf{f}_{Y,0}(\cdot|A_n, W_n)\mu_Y$, respectively, and (iii) that each A_i is drawn conditionally on W_i from the Bernoulli distribution with known parameter $g_i(1|W_i)$.

Let \mathcal{F} be the semiparametric collection of all elements of the form

$$\mathbf{f} = (\mathbf{f}_W, \mathbf{f}_Y(\cdot|a, w), (a, w) \in \mathscr{A} \times \mathscr{W})$$

where \mathbf{f}_W is a density wrt μ_W and each $\mathbf{f}_Y(\cdot|a, w)$ is a density wrt μ_Y . In particular, we define $\mathbf{f}_0 \equiv (\mathbf{f}_{W,0}, \mathbf{f}_{Y,0}(\cdot|a, w), (a, w) \in \mathscr{A} \times \mathscr{W}) \in \mathcal{F}$. In light of (2.1) define, for every $\mathbf{f} \in \mathcal{F}$, $\mathcal{L}_{\mathbf{f}, \mathbf{g}_n}(\mathbf{O}_n) \equiv \prod_{i=1}^n \mathbf{f}_W(W_i) \times g_i(A_i|W_i) \times \mathbf{f}_Y(Y_i|A_i, W_i)$. The set $\{\mathcal{L}_{\mathbf{f}, \mathbf{g}_n} : \mathbf{f} \in \mathcal{F}\}$ is a semiparametric model for the likelihood of \mathbf{O}_n . It contains the true, unknown likelihood $\mathcal{L}_{\mathbf{f}_0, \mathbf{g}_n}$.

For the sake of illustration, we choose the marginal treatment effect on an additive scale as our parameter of interest. Thus, let $\Upsilon : \mathcal{F} \rightarrow [-1, 1]$ be the mapping such that, for every $\mathbf{f} = (\mathbf{f}_W, \mathbf{f}_Y(\cdot|a, w), (a, w) \in \mathscr{A} \times \mathscr{W})$,

$$\Upsilon(\mathbf{f}) = \int (Q_{Y,\mathbf{f}}(1, w) - Q_{Y,\mathbf{f}}(0, w)) \mathbf{f}_W(w) d\mu_W, \quad (2.2)$$

where $Q_{Y,\mathbf{f}}(a, w) = \int y \mathbf{f}_Y(y|a, w) d\mu_Y$ is the mean of $\mathbf{f}_Y(\cdot|a, w)\mu_Y$. The true marginal treatment effect on an additive scale is $\psi_0 \equiv \Upsilon(\mathbf{f}_0)$. Of particular interest in medical, epidemiological and social sciences research, it can be interpreted causally under assumptions on the data-generating process (Pearl, 2000).

We have not specified yet what is precisely the sequence of randomization schemes $\mathbf{g}_n \equiv (g_1, \dots, g_n)$. Our CARA sampling scheme “targets” a randomization scheme g_0 which minimizes a user-supplied optimality criterion. By targeting g_0 we mean estimating

g_0 based on past observations, and relying on the resulting estimator to collect the next block of data, as seen in (2.1). For the sake of illustration, we consider the case that g_0 is the following minimizer

$$g_0 \equiv \arg \min_g E_{P_{Q_0, g}} \left(\frac{(Y - Q_{Y, \mathbf{f}_0}(A, W))^2}{g^2(A|W)} \right) \quad (2.3)$$

across all randomization schemes g . We emphasize that the above definition of g_0 involves the unknown \mathbf{f}_0 , so it is unknown too. We will comment on (2.3) and motivate our interest in g_0 in Section 2.3. As we shall see, g_0 minimizes a generalized Cramér-Rao lower bound for ψ_0 . Known in the literature as the *Neyman design* (Hu and Rosenberger, 2006), g_0 actually has a closed-form expression as a function of \mathbf{f}_0 . We *do not* use this closed-form expression in order to illustrate the generality of our framework which allows to target any randomization scheme defined as a minimizer of an optimality criterion.

Notation

Let $O \equiv (W, A, Y)$ be a generic data-structure. Every distribution of O consists of two components: on the one hand, the marginal distribution of W and the conditional distribution of Y given (A, W) , which correspond to a $\mathbf{f} \in \mathcal{F}$; on the other hand, the conditional distribution of A given W , or randomization scheme. To reflect this dichotomy, we denote the distribution of O as $P_{Q, g}$, where Q equals the couple formed by the marginal distribution of W and the conditional distribution of Y given (A, W) , and g equals the randomization scheme. We denote Q_0 the true couple Q in our population of interest, which corresponds to \mathbf{f}_0 and is unknown to us. For a given Q , we denote Q_W the related marginal distribution of W and Q_Y the related conditional expectation of Y given (A, W) . If $Q = Q_0$, then Q_W and Q_Y are denoted $Q_{W,0}$ and $Q_{Y,0}$, respectively.

We denote \mathcal{G} and \mathcal{Q}_Y the set of all randomization schemes and the set of all conditional expectations of Y given (A, W) , respectively. Thus, for any Q and g , $P_{Q_0, g}$ is the true, partially unknown distribution of O when one relies on g , and $E_{P_{Q_0, g}}(Y|A, W) = Q_Y(A, W)$, $P_{Q_0, g}(A = 1|W) = g(1|W) = 1 - g(0|W)$ $P_{Q_0, g}$ -almost surely.

With this notation, ψ_0 can be rewritten

$$\psi_0 = \int (Q_{Y,0}(1, w) - Q_{Y,0}(0, w)) dQ_{W,0}(w)$$

and satisfies $\psi_0 = E_{P_{Q_0, g}}(Q_{Y,0}(1, W) - Q_{Y,0}(0, W))$ whatever is $g \in \mathcal{G}$.

Loss functions and working models

Let $g^b \in \mathcal{G}$ be the balanced randomization scheme wherein each arm is assigned with probability $1/2$ regardless of baseline covariates. Let $g^r \in \mathcal{G}$ be a randomization scheme, bounded away from 0 and 1 by choice, that serves as a reference. This can be simply be the balanced scheme. In addition, let L be a loss function for $Q_{Y,0}$ and $\mathcal{Q}_{1,n}$ be a working model for the conditional response

$$\mathcal{Q}_{1,n} \equiv \{Q_{Y,\beta} : \beta \in B_n\} \subset \mathcal{Q}_Y.$$

These working models can be a series of LASSO models that increase with sample size n . One choice of L is the quasi negative-log-likelihood loss function L^{kl} . For any $Q_Y \in \mathcal{Q}_Y$ bounded away from 0 and 1, $L^{\text{kl}}(Q_Y)$ satisfies

$$-L^{\text{kl}}(Q_Y)(O) \equiv Y \log(Q_Y(A, W)) + (1 - Y) \log(1 - Q_Y(A, W)). \quad (2.4)$$

Another interesting loss function L for $Q_{Y,0}$ is the least-square loss function L^{ls} , given by

$$L^{\text{ls}}(Q_Y)(O) \equiv (Y - Q_Y(A, W))^2. \quad (2.5)$$

Likewise, let L_{Q_Y} be a loss function for g_0 , which may depend on $Q_Y \in \mathcal{Q}_Y$, and let $\mathcal{G}_{1,n} \subset \mathcal{G}$ be a working model for the optimal randomization scheme. These can be a series of LASSO models that increase with sample size n . In this context, a loss function for g_0 may be given, for any $Q_Y \in \mathcal{Q}_Y$, by

$$L_{Q_Y}(g)(O) \equiv \frac{(Y - Q_Y(A, W))^2}{g(A|W)}. \quad (2.6)$$

We explain the motivation and justification for this loss function in section 2.3.

As suggested by the notation, the sets $\mathcal{Q}_{1,n}$ and $\mathcal{G}_{1,n}$ may depend on n . In that case, the two sequences of sets must be non-decreasing. Moreover, the specifications must guarantee that the elements of $\mathcal{Q}_1 \equiv \cup_{n \geq 1} \mathcal{Q}_{1,n}$ and those of $\mathcal{G}_1 \equiv \cup_{n \geq 1} \mathcal{G}_{1,n}$ be uniformly bounded away from 0 and 1.

Targeted adaptive sampling and inference

The estimation of g_0 involves the estimation of $Q_{Y,0}$. At each step, the current estimators of $Q_{Y,0}$ and g_0 are also used to craft a targeted estimator of ψ_0 .

We initialize the sampling scheme by setting $g_1 \equiv g^b$. Consider $1 < i < n$. Since

$$g_0 = \arg \min_{g \in \mathcal{G}} E_{P_{Q_0, s}} \left(\frac{L_{Q_{Y,0}}(g)(O)}{g(A|W)} \right) \quad \text{and} \quad Q_{Y,0} = \arg \min_{Q_Y \in \mathcal{Q}} E_{P_{Q_0, s}} (L(Q_Y)(O)),$$

we define

$$g_i \in \arg \min_{g \in \mathcal{G}_{1,i}} \frac{1}{i-1} \sum_{j=1}^{i-1} \frac{L_{Q_{Y,\beta_i}}(g)(O_j)}{g_j(A_j|W_j)}, \quad (2.7)$$

where

$$\beta_i \in \arg \min_{\beta \in \mathcal{B}_i} \frac{1}{i-1} \sum_{j=1}^{i-1} L(Q_{Y,\beta})(O_j) \frac{g^r(A_j|W_j)}{g_j(A_j|W_j)}. \quad (2.8)$$

These weights will play an important role in establishing the convergence of the randomization schemes. The reference scheme in the numerator is necessary to obtain the convergence of the outcome models before establishing the convergence of the randomization schemes. By specifying the sequence of randomization schemes, this completes the definition of the likelihood function, hence the characterization of our sampling scheme.

To estimate ψ_0 based on \mathbf{O}_n , we introduce the following one-dimensional parametric model for $Q_{Y,0}$:

$$\{Q_{Y,\beta_n}(\varepsilon) \equiv \text{expit}(\text{logit}(Q_{Y,\beta_n}) + \varepsilon H(g_n)) : \varepsilon \in \mathcal{E}\}, \quad (2.9)$$

where $\mathcal{E} \subset \mathbb{R}$ is a closed, bounded interval containing 0 in its interior and $H(g)(O) \equiv (2A-1)/g(A|W)$. The optimal fluctuation parameter is one that minimizes an empirical risk

$$\varepsilon_n \in \arg \min_{\varepsilon \in \mathcal{E}} \frac{1}{n} \sum_{i=1}^n L^{\text{kl}}(Q_{Y,\beta_n}(\varepsilon))(O_i) \frac{g_n(A_i|W_i)}{g_i(A_i|W_i)}. \quad (2.10)$$

We set $Q_{Y,\beta_n}^* \equiv Q_{Y,\beta_n}(\varepsilon_n)$ and define

$$\psi_n^* \equiv \frac{1}{n} \sum_{i=1}^n Q_{Y,\beta_n}^*(1, W_i) - Q_{Y,\beta_n}^*(0, W_i). \quad (2.11)$$

We show in Section 2.4 that ψ_n^* consistently estimates ψ_0 , and that $\sqrt{n}/\Sigma_n(\psi_n^* - \psi_0)$ is approximately standard normally distributed, where Σ_n is an explicit estimator (2.20). This enables the construction of confidence intervals of desired asymptotic level. As for the optimal randomization scheme g_0 , we show that it is targeted indeed, in the sense that g_n converges to the projection of g_0 onto $\cup_{n \geq 1} \mathcal{G}_{1,n}$. The assumptions under which our results hold are typical of loss-based inference. They essentially concern the existence and convergence of projections, as well as the complexity of our working models, expressed in terms of bracketing numbers and integrals. In Section 2.5, we develop and study a specific example based on the LASSO.

Further notation

Consider a class \mathcal{F} of real-valued functions and $\phi : \mathbb{R} \rightarrow \mathbb{R}$. If $\phi \circ f$ is well-defined for each $f \in \mathcal{F}$, then we note $\phi(\mathcal{F}) \equiv \{\phi \circ f : f \in \mathcal{F}\}$. Set a semi-norm $\|\cdot\|$ on \mathcal{F} and $\delta > 0$. We denote $N(\delta, \mathcal{F}, \|\cdot\|)$ the δ -bracketing number of \mathcal{F} wrt $\|\cdot\|$ and $J(\delta, \mathcal{F}, \|\cdot\|) \equiv \int_0^\delta \sqrt{\log N(\varepsilon, \mathcal{F}, \|\cdot\|)} d\varepsilon$ the corresponding bracketing integral at δ .

In general, given a known $g \in \mathcal{G}$ and an observation O drawn from $P_{Q_0, g}$, $Z \equiv g(A|W)$ is a deterministic function of g and O . Note that Z should be interpreted as a weight associated with O and will be used as such. Therefore, we can augment O with Z , *i.e.*, substitute (O, Z) for O , while still denoting $(O, Z) \sim P_{Q_0, g}$. In particular, during the course of our trial, conditionally on \mathbf{O}_{i-1} , the randomization scheme g_i is known and we can substitute $(O_i, Z_i) = (O_i, g_i(A_i|W_i)) \sim P_{Q_0, g_i}$ for O_i drawn from P_{Q_0, g_i} . The inverse weights $1/g_i(A_i|W_i)$ are bounded because \mathcal{G}_1 is uniformly bounded away from 0 and 1.

The empirical distribution of \mathbf{O}_n is denoted P_n . For a function $f : \mathcal{O} \times [0, 1] \rightarrow \mathbb{R}^d$, we use the notation $P_n f \equiv n^{-1} \sum_{i=1}^n f(O_i, Z_i)$. Likewise, for any fixed $P_{Q, g} \in \mathcal{M}$, $P_{Q, g} f \equiv E_{P_{Q, g}}(f(O, Z))$ and, for each $i = 1, \dots, n$, $P_{Q_0, g_i} f \equiv E_{P_{Q_0, g_i}}[f(O_i, Z_i)|\mathbf{O}_{i-1}]$, and $P_{Q_0, \mathbf{g}_n} f \equiv n^{-1} \sum_{i=1}^n E_{P_{Q_0, g_i}}[f(O_i, Z_i)|\mathbf{O}_{i-1}]$.

The supremum norm of a function $f : \mathcal{O} \times [0, 1] \rightarrow \mathbb{R}^d$ is denoted $\|f\|_\infty$. If $d = 1$ and f is measurable, then the $L^2(P_{Q_0, g^r})$ -norm of f is given by $\|f\|_{2, P_{Q_0, g^r}}^2 \equiv P_{Q_0, g^r} f^2$. If f is only a function of W , then we denote $\|f\|_{2, Q_{W, 0}}$ its $L^2(P_{Q_0, g^r})$ -norm, to emphasize that it only depends on the marginal distribution $Q_{W, 0}$. With a slight abuse of notation, if f is only a function of (A, W) , then $\|f\|_{2, Q_{W, 0}}^2$ is the $L^2(P_{Q_0, g^r})$ -norm of $w \mapsto f(1, w)$. In particular, for $Q_Y, Q'_Y \in \mathcal{Q}_Y$ and $g, g' \in \mathcal{G}$, $\|Q_Y - Q'_Y\|_{2, P_{Q_0, g^r}}^2 = E_{P_{Q_0, g^r}}(Q_Y(A, W) - Q'_Y(A, W))^2$, and $\|g - g'\|_{2, Q_{W, 0}}^2 = E_{Q_{W, 0}}((g(1|W) - g'(1|W))^2)$.

Asymptotics

Our main result rely on the following assumptions.

A1. The conditional distribution of Y given (A, W) under Q_0 is not degenerated.

Existence and convergence of projections.

A2 For each $n \geq 1$, there exists $Q_{Y, \beta_{n, 0}} \in \mathcal{Q}_{1, n}$ satisfying

$$P_{Q_0, g^r} L(Q_{Y, \beta_{n, 0}}) = \inf_{Q_{Y, \beta} \in \mathcal{Q}_{1, n}} P_{Q_0, g^r} L(Q_{Y, \beta}).$$

There also exists $Q_{Y,\beta_0} \in \mathcal{Q}_1$ such that, for all $\delta > 0$,

$$P_{Q_0, g^r} L(Q_{Y,\beta_0}) < \inf_{\{Q_Y \in \mathcal{Q}_1: \|Q_Y - Q_{Y,\beta_0}\|_{2, P_{Q_0, g^r}} \geq \delta\}} P_{Q_0, g^r} L(Q_Y).$$

A3. For each $n \geq 1$, there exists $g_{n,0} \in \mathcal{G}_{1,n}$ satisfying

$$P_{Q_0, g^r} L_{Q_{Y,\beta_0}}(g_{n,0})/g^r = \inf_{g \in \mathcal{G}_{1,n}} P_{Q_0, g^r} L_{Q_{Y,\beta_0}}(g)/g^r. \quad (2.12)$$

There also exists $g_0^* \in \mathcal{G}_1$ such that, for all $\delta > 0$,

$$P_{Q_0, g^r} L_{Q_{Y,\beta_0}}(g_0^*)/g^r < \inf_{\{g \in \mathcal{G}_1: \|g - g_0^*\|_{2, Q_{W,0}} \geq \delta\}} P_{Q_0, g^r} L_{Q_{Y,\beta_0}}(g)/g^r. \quad (2.13)$$

A4. Assume that Q_{Y,β_0} from **A2** and g_0^* from **A3** exist. For each $\varepsilon \in \mathcal{E}$, introduce

$$Q_{Y,\beta_0}(\varepsilon) \equiv \text{expit}(\text{logit}(Q_{Y,\beta_0}) + \varepsilon H(g_0^*)), \quad (2.14)$$

where $H(g_0^*)(O) \equiv (2A - 1)/g_0^*(A|W)$. Then, there is a unique $\varepsilon_0 \in \mathcal{E}$ such that

$$\varepsilon_0 \in \arg \min_{\varepsilon \in \mathcal{E}} P_{Q_0, g_0^*} L^{\text{kl}}(Q_{Y,\beta_0}(\varepsilon)). \quad (2.15)$$

Reasoned complexity.

A5. $J(1, \mathcal{Q}_{1,n}, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(\sqrt{n})$ and $J(1, L(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, g^r}}) = o(\sqrt{n})$.

A5*. If $\{\delta_n\}_{n \geq 1}$ is a sequence of positive numbers, then $\delta_n = o(1)$ implies $J(\delta_n, \mathcal{Q}_{1,n}, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$.

A6. The entropy condition $J(1, \mathcal{G}_{1,n}, \|\cdot\|_{2, Q_{W,0}}) = o(\sqrt{n})$ holds.

A6*. If $\{\delta_n\}_{n \geq 1}$ is a sequence of positive numbers, then $\delta_n = o(1)$ implies $J(\delta_n, \mathcal{G}_n, \|\cdot\|_{2, Q_{W,0}}) = o(1)$.

Theorem 2.1 (Asymptotic study of the targeted CARA RCT). *Assume that **A2**, **A3**, **A5** and **A6** are met. Then, the targeted CARA design converges in the sense that $\|g_n - g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ in probability as $n \rightarrow \infty$. If, in addition, **A4** holds, then the TMLE ψ_n^* consistently estimates ψ_0 . Moreover, if **A1**, **A5*** and **A6*** also hold, then $\sqrt{n}/\Sigma_n(\psi_n^* - \psi_0)$ is approximately standard normally distributed, where Σ_n is the explicit estimator given in (2.20).*

The last statement in the above theorem underpins the statistical analysis of the proposed targeted CARA RCT. In particular, denoting $\xi_{1-\alpha/2}$ the $(1 - \alpha/2)$ -quantile of the standard normal distribution, the interval $[\psi_n^* \pm \xi_{1-\alpha/2} \sqrt{\Sigma_n/n}]$ is a confidence interval of asymptotic level $(1 - \alpha)$.

2.3 Comments on Section 2.2

A closer look at the parameter of interest and the optimal randomization scheme

Central to our approach is formulating ψ_0 as the value at \mathbf{f}_0 of the mapping $\Upsilon : \mathcal{F} \rightarrow [-1, 1]$ given by (2.2). Let \mathcal{M} denote the set of all possible distributions of O . Because we slightly changed perspective and now think in terms of distributions $P_{Q,g} \in \mathcal{M}$ instead of $\mathbf{f} \in \mathcal{F}$, it is convenient to introduce the mapping $\Psi : \mathcal{M} \rightarrow [-1, 1]$ characterized by

$$\Psi(P_{Q,g}) \equiv \int (Q_Y(1, w) - Q_Y(0, w)) dQ_W(w) = E_{P_{Q,g}}(Q_Y(1, W) - Q_Y(0, W)).$$

Since Ψ only depends on $P_{Q,g}$ through Q , we will now on systematically write $\Psi(Q)$ in place of $\Psi(P_{Q,g})$ to alleviate notation.

The mapping Ψ is pathwise differentiable. Its efficient influence curve sheds light on the asymptotic properties of all regular and asymptotically linear estimators of $\psi_0 = \Psi(Q_0)$. The latter statement is formalized in the following lemma—we refer the reader to (Bickel, Klaassen, Ritov, and Wellner, 1998, van der Laan and Robins, 2003, van der Vaart, 1998) for definitions and proofs.

Lemma 2.1. *The mapping $\Psi : \mathcal{M} \rightarrow [-1, 1]$ is pathwise differentiable at every $P_{Q,g} \in \mathcal{M}$ wrt the maximal tangent space. Its efficient influence curve at $P_{Q,g}$, denoted $D(P_{Q,g})$, orthogonally decomposes as $D(P_{Q,g})(O) = D_W(Q)(W) + D_Y(Q_Y, g)(O)$ with*

$$\begin{aligned} D_W(Q)(W) &\equiv Q_Y(1, W) - Q_Y(0, W) - \Psi(Q), \\ D_Y(Q_Y, g)(O) &\equiv \frac{2A - 1}{g(A|W)} (Y - Q_Y(A, W)). \end{aligned}$$

The variance $\text{Var}_{P_{Q,g}}(D(P)(O))$ is a generalized Cramér-Rao lower bound for the asymptotic variance of any regular and asymptotically linear estimator of $\Psi(Q)$ when sampling independently from $P_{Q,g}$.

Moreover, if either $Q_Y = Q'_Y$ or $g = g'$ then $E_{P_{Q,g}}(D(P_{Q',g'})(O)) = 0$ implies $\Psi(Q) = \Psi(Q')$.

The last statement of Lemma 2.1, often referred to as a “double-robustness” property, assures that D can be deployed to safeguard against model mis-specifications when estimating ψ_0 . This is especially relevant in an RCT setting, since the randomization scheme g is known whenever one samples an observation from $P_{Q,g}$.

By Lemma 2.1, the asymptotic variance of a regular, asymptotically linear estimator under independent sampling from $P_{Q_0, g}$ is lower-bounded by

$$\min_{g \in \mathcal{G}} \text{Var}_{P_{Q_0, g}}(D(P_{Q_0, g})(O)) = \min_{g \in \mathcal{G}} E_{P_{Q_0, g}} \left(\frac{(Y - Q_{Y,0}(A, W))^2}{g^2(A|W)} \right).$$

In this light, targeting g_0 defined by (2.3) means that the goal of adaptation is to reach a randomization scheme of higher efficiency, *i.e.*, to obtain a valid estimate of ψ_0 using as few blocks of patients as possible. As mentioned in section 2.2, though not used in our approach, g_0 actually has a closed form expression $g_0(1|W) = \sigma_0(1, W) / (\sigma_0(1, W) + \sigma_0(0, W))$, where $\sigma_0^2(A, W)$ is the conditional variance of Y given (A, W) under Q_0 . Under this randomization scheme, the treatment arm with higher probability for a patient with baseline covariates W is the one for which the conditional variance of the response is higher.

On the data-adaptive loss-based estimation of $Q_{Y,0}$

The reference randomization scheme g^r offers the opportunity to differentially weight each observation in (2.8). This action impacts the convergence of Q_{Y, β_n} and thus that of g_n , as seen in Sections 2.2 and 2.2 (the limit g_0^* depends on g^r).

As we already emphasized, the working model $\mathcal{Q}_{1, n}$ may depend on sample size n . If it does, then the sequence of working models must be non-decreasing and $\mathcal{Q}_1 \equiv \cup_{n \geq 1} \mathcal{Q}_{1, n}$ can be interpreted as the limiting working model for $Q_{Y,0}$. We would typically recommend to start with $\mathcal{Q}_1 = \dots = \mathcal{Q}_{n_0}$ all equal to a small set, with a user-supplied, deterministic n_0 , then to let the complexity grow with n . It is known, however, that such a growth must remain tethered. Assumptions **A5** and **A5*** provide appropriate conditions on the complexity of $\mathcal{Q}_{1, n}$. We refer to Section 2.3 for a discussion of their meaning.

The combined choice of loss function L and working model $\mathcal{Q}_{1, n}$ determines the technique used to estimate $Q_{Y,0}$. For instance, in the traditional parametric approach, the working model $\mathcal{Q}_{1, n}$ does not depend on n and is indexed by a fixed, finite-dimensional parameter set. Under the LASSO methodology, which we carefully describe and study in Section 2.5, $\text{logit}(\mathcal{Q}_{1, n})$ is the linear span of a given basis, with constraints on the linear combinations imposed through the definition of B_n .

On the data-adaptive loss-based estimation of g_0

The optimal randomization scheme g_0 is defined as a minimizer of a certain criterion over the class \mathcal{G} of all randomization schemes, see (2.3). Thus, our loss-based estimation of g_0 based on \mathbf{O}_n consists of defining g_{n+1} as the minimizer in g of an estimator of the

optimality criterion over the user-supplied class of randomization schemes $\mathcal{G}_{1,n}$, see (2.7) and the next paragraph. This approach is applicable in the largest generality. Alternatively, if W is discrete, then g_0 takes finitely many values and g_{n+1} can be defined explicitly based on Q_{Y,β_n} and \mathbf{O}_n . This is also the case if one is willing to assign treatment only based on a discrete summary measure V of W . In this context, g_0 is defined as in (2.3), where the arg min is over the subset of \mathcal{G} consisting of those randomization schemes which depend on W only through V . We refer the readers to (Chambaz and van der Laan, 2013) for details. Note that assigning treatment based on such summary measures is perhaps too restrictive in real-life RCTs where response to treatment may be correlated with a large number of a patient's baseline characteristics, some of which being continuous.

We now turn to the joint justification of (2.6) and (2.7). The key point is the following equality, valid for every $g' \in \mathcal{G}$:

$$g_0 = \arg \min_{g \in \mathcal{G}} E_{P_{Q_0, g'}} \left(\frac{(Y - Q_{Y,0}(A, W))^2}{g(A|W)g'(A|W)} \right). \quad (2.16)$$

Equality (2.16) tells us that g_0 can be estimated using observations drawn from $P_{Q_0, g'}$ based on the loss function L_{Q_Y} provided it is weighted by $1/g'$. Our observations O_1, \dots, O_n are drawn from $P_{Q_0, g_1}, \dots, P_{Q_0, g_n}$, respectively. In this light, (2.16) also validates (2.7). At each step of our modification, we are given a data-generating scheme g' , and if we were also given $Q_{Y,0}$, what would be the optimal randomization scheme. Therefore, an estimator (subsequent randomization) that targets this g_0 using observations generated by g_1, \dots, g_n would give us a step closer to g_0 .

Like $\mathcal{D}_{1,n}$, the working model $\mathcal{G}_{1,n} \subset \mathcal{G}$ may depend on sample size n . If it does, then the sequence must be non-decreasing and $\mathcal{G}_1 \equiv \cup_{n \geq 1} \mathcal{G}_{1,n}$ can be interpreted as the limiting working model for g_0 . The additional constraint that \mathcal{G}_1 be uniformly bounded away from 0 and 1 is important. It implies the following pivotal property: no matter how $g^r \in \mathcal{G}$ is chosen in Section 2.2, there exists some constant $\kappa > 0$, such that $\max(\|g^r/g\|_\infty, \|g/g^r\|_\infty) \leq \kappa$, for all $g \in \mathcal{G}_1$. For ease of reference, we call it the *dominated ratio property* of \mathcal{G}_1 .

Using a fixed working model for g_0 , *i.e.*, setting $\mathcal{G}_1 = \mathcal{G}_{1,n}$ for every $n \geq 1$, is a valuable option. However, in some situations, *e.g.* if the population is very heterogeneous, using a fixed, large working model \mathcal{G}_1 may delay, sample size-wise, the adaptation, thereby depriving the trial of the advantages of an adaptive design. By allowing $\mathcal{G}_{1,n}$ to depend on n , one gains the flexibility to enrich the working model for g_0 according to the modesty or generosity of the sample size. Similar to what we suggested for $\mathcal{D}_{1,n}$ in Section 2.3, we would recommend to start with $\mathcal{G}_1 = \dots = \mathcal{G}_{n_0}$ all equal to a small set, typically the singleton $\{g^b\}$, then to let the complexity of $\mathcal{G}_{1,n}$ augment with n , though not too abruptly. Assumptions **A6** and **A6*** provide appropriate conditions on the complexity of $\mathcal{G}_{1,n}$. We

refer to Section 2.3 for a discussion of their meaning. The assumptions are mild, and allow us to use the LASSO to target the optimal randomization scheme g_0 , just like we can use the LASSO to estimate $Q_{Y,0}$, see Section 2.5.

On targeted minimum loss estimation

The conception of ψ_n^* defined in (2.11) follows the paradigm of targeted minimum loss estimation. In the setting of a covariate-adjusted RCT with a fixed design and a fixed working model \mathcal{Q}_1 , a TMLE estimator is unbiased and asymptotically Gaussian regardless of the specification of \mathcal{Q}_1 . Chambaz and van der Laan (2013) show that unbiasedness and asymptotic normality still hold in a framework very similar to that of the present article when the randomization schemes depend on W only through a summary measure taking finitely many values and when \mathcal{Q}_1 is a simple linear model. Such a configuration can be obtained as a particular case of the example developed in Section 2.5.

Although using a mis-specified parametric working model \mathcal{Q}_1 for $Q_{Y,0}$ does not hinder the consistency of the estimator of ψ_0 , it may affect its efficiency and the convergence of the CARA design to the targeted optimal design. By relying on more flexible randomization schemes and on more adaptive estimators of $Q_{Y,0}$, we may better adapt to the optimal randomization scheme g_0 through better variable adjustments and the targeted construction of the instrumental loss function L_{Q_Y} . Because g_0 is the Neyman design, our approach yields greater efficiency through better variable adjustments and more accurate estimation of the variance of the estimator.

Consider now (2.9) and (2.11). The model (2.9) goes through Q_{Y,β_n} at $\varepsilon = 0$ and satisfies the score condition $\frac{\partial}{\partial \varepsilon} L^{\text{kl}}(Q_{Y,\beta_n}(\varepsilon))|_{\varepsilon=0} = D_Y(Q_{Y,\beta_n}, g_n)$. If we set $Q_{\beta_n}^* \equiv (Q_{W,n}, Q_{Y,\beta_n}^*)$, where $Q_{W,n}$ is the empirical marginal distribution of W , then $\psi_n^* = \Psi(Q_{\beta_n}^*)$, assuring that ψ_n^* is indeed a substitution estimator of $\psi_0 = \Psi(Q_0)$. The use of substitution principle allows one to preserve global information embedded in the parameter map, such as the bounds of the parameter space, and this may provide further finite sample gain, though would not make a difference in asymptotic behavior.

On the assumptions

Assumption **A2** stipulates the existence of a projection $Q_{Y,\beta_{n,0}}$ of $Q_{Y,0}$ onto every working model $\mathcal{Q}_{1,n}$. In other words, on every working model, there is a limiting function minimizing the true risk within this model. In particular, one does not assume that these working models converge to the unknown true response model $Q_{Y,0}$. It may depend on the user-supplied reference randomization scheme g^r . If $Q_{Y,0} \in \mathcal{Q}_1$, *i.e.*, if \mathcal{Q}_0 is well-specified, then the existence of $Q_{Y,\beta_0} = Q_{Y,0}$ is granted. If $Q_{Y,0} \notin \mathcal{Q}_1$, *i.e.*, if \mathcal{Q}_1 is mis-specified,

then **A2** also stipulates the existence of a projection Q_{Y,β_0} of $Q_{Y,0}$ onto \mathcal{Q}_1 . It may also depend on g^r .

Similar comments apply to **A3**. Note that each $g_{n,0}$ and the limiting randomization scheme g_0^* depend on g^r only through Q_{Y,β_0} : replacing g^r with any arbitrarily chosen $g \in \mathcal{G}$ in (2.12) or (2.13) does not alter the values of $g_{n,0}$ and g_0^* . Furthermore, (2.6) and (2.13) yield that

$$g_0^* = \arg \min_{g \in \mathcal{G}_1} \left\{ \text{Var}_{P_{Q_0,g}}(D_Y(Q_{Y,0},g)(O)) + P_{Q_0,g} \frac{(Q_{Y,0} - Q_{Y,\beta_0})^2}{g^2} \right\}.$$

This shows that if $Q_{Y,\beta_0} = Q_{Y,0}$ and $g_0 \in \mathcal{G}_1$, then $g_0^* = g_0$, the optimal randomization scheme. In general, g_0^* minimizes an objective function which is the sum of the Cramér-Rao lower bound and a second-order residual. This underscores the motivation for using a flexible estimator in estimating $Q_{Y,0}$: by minimizing the second-order residual, we get closer to adapting towards the desired optimal randomization criterion.

Recall that Q_{Y,β_n} is characterized by (2.8) and that $Q_{W,n}$ is the empirical marginal distribution of W . Heuristically, if the equality $Q_{Y,\beta_0} = Q_{Y,0}$ holds then one should be able to prove that $\Psi((Q_{W,n}, Q_{Y,\beta_n}))$ is a consistent estimator of ψ_0 . Since $Q_{Y,\beta_0} = Q_{Y,0}$ also yields that $\varepsilon_0 = 0$ is the unique solution to (2.15) in **A5**, one understands that updating Q_{Y,β_n} to $Q_{Y,\beta_n}^* \equiv Q_{Y,\beta_n}(\varepsilon_n)$ and $\Psi((Q_{W,n}, Q_{Y,\beta_n}))$ to ψ_n^* as described in (2.10) and (2.11) should preserve the consistency in the initially well-specified framework. In the more likely situation where \mathcal{Q}_1 is mis-specified, hence $Q_{Y,\beta_0} \neq Q_{Y,0}$ and $\varepsilon_0 \neq 0$, there is no reason to believe that $\Psi((Q_{W,n}, Q_{Y,\beta_n}))$ should be a consistent estimator of ψ_0 . In this light, the updating procedure bends the inconsistent initial estimator into a consistent one by drawing advantage from the double-robustness of D that we presented in Lemma 2.1.

In Sections 2.3 and 2.3, we commented on the interest of letting the working models $\mathcal{Q}_{1,n}$ and $\mathcal{G}_{1,n}$ depend on sample size n . Assumptions **A5** and **A6** put very mild constraints on how the complexities of the working models should evolve with n to guarantee the convergence of g_n and consistency of ψ_n^* . The constraints are expressed in terms of bracketing integral. We refer the reader to (van der Vaart, 1998, Examples 19.7-19.11, Lemma 19.15) for typical examples. They include “well-behaved” parametric and Vapnik-Cervonenkis (VC) classes. Assumptions **A5*** and **A6*** should be interpreted as more stringent conditions imposed upon $\mathcal{Q}_{1,n}$ and $\mathcal{G}_{1,n}$. Indeed, for instance,

$$J(1, \mathcal{G}_{1,n}, \|\cdot\|_{2, Q_{W,0}}) / \sqrt{n} \leq J(1/\sqrt{n}, \mathcal{G}_{1,n}, \|\cdot\|_{2, Q_{W,0}})$$

because the entropy with bracketing is non-increasing, so that **A6*** does imply **A6** (take $\delta_n = 1/\sqrt{n}$). The need for more stringent conditions arises when studying the convergence in law of ψ_n^* .

2.4 Building blocks of Theorem 2.1

We now carry out the theoretical study of the targeted CARA design and its corresponding estimator described in Section 2.2. All proofs are relegated to Section 2.8.

We first focus on the convergence of the estimators Q_{Y,β_n} . The counterpart to this result in the i.i.d. setting is well established (Pollard, 1984, van der Vaart, 1998, among others). The following proposition revises those results for the current statistical setting.

Proposition 2.1 (convergence of Q_{Y,β_n}). *Under A2, A5, $\|Q_{Y,\beta_n} - Q_{Y,\beta_0}\|_{2,P_{Q_0,g^r}} \rightarrow 0$ in probability as $n \rightarrow \infty$.*

We now turn to the convergence of the sequence of randomization schemes.

Proposition 2.2 (convergence of the targeted CARA design). *Under A2, A3, A5 and A6, it holds that $\|g_n - g_0^*\|_{2,Q_{W,0}} \rightarrow 0$ in probability as $n \rightarrow \infty$.*

The following corollary of Proposition 2.2 will also prove useful.

Corollary 2.1. *Assume the setting of Proposition 2.2.*

It also holds that $\|g_n - g_0^\|_{2,Q_{W,0}} \rightarrow 0$ in L^1 as $n \rightarrow \infty$, and that $\|1/g_n - 1/g_0^*\|_{2,Q_{W,0}}$, $\|n^{-1} \sum_{i=1}^n g_i - g_0^*\|_{2,Q_{W,0}}$, and $\|n^{-1} \sum_{i=1}^n 1/g_i - 1/g_0^*\|_{2,Q_{W,0}}$ converge to 0 in probability and in L^1 as $n \rightarrow \infty$.*

At this stage, the consistency of ψ_n^* can be established. The proof relies on the convergence of Q_{Y,β_n}^* to a limiting conditional distribution, which is a fluctuation of the limit Q_{Y,β_0} of Q_{Y,β_n} , see Proposition 2.1.

Proposition 2.3 (consistency of ψ_n^*). *Suppose that A2, A3, A4, A5 and A6 are met. Define*

$$Q_{Y,\beta_0}^* \equiv \text{expit}(\text{logit}(Q_{Y,\beta_0}) + \varepsilon_0 H(g_0^*)), \quad (2.17)$$

with $H(g_0^)(O) \equiv (2A - 1)/g_0^*(A|W)$ and $Q_{\beta_0}^* \equiv (Q_{W,0}, Q_{Y,\beta_0}^*)$. It holds that $\|Q_{Y,\beta_n}^* - Q_{Y,\beta_0}^*\|_{2,P_{Q_0,g^r}} \rightarrow 0$ in probability as $n \rightarrow \infty$. Moreover, $\Psi(Q_{\beta_0}^*) = \psi_0$ and ψ_n^* consistently estimates ψ_0 .*

We need further notation to state our last building block. For both $\beta = \beta_0$ and $\beta = \beta_n$, introduce $d_{Y,\beta}^*$ given by

$$d_{Y,\beta}^*(O, Z) \equiv \frac{2A - 1}{Z} \left(Y - Q_{Y,\beta}^*(A, W) \right). \quad (2.18)$$

Define also

$$\Sigma_0 \equiv P_{Q_0, g_0^*} \left(d_{Y, \beta_0}^* + D_W(Q_{\beta_0}^*) \right)^2 = P_{Q_0, g_0^*} \left(D(P_{Q_{\beta_0}^*, g_0^*}) \right)^2, \quad (2.19)$$

$$\Sigma_n \equiv \frac{1}{n} \sum_{i=1}^n \left(d_{Y, \beta_n}^*(O_i, Z_i) + D_W(Q_{\beta_n}^*)(W_i) \right)^2, \quad (2.20)$$

where we recall that $Q_{\beta_n}^* \equiv (Q_{W, n}, Q_{Y, \beta_n}^*)$.

Proposition 2.4 (asymptotic linearity and central limit theorem for ψ_n^*). *Assume that AI–A6* are met. Then $\Sigma_n = \Sigma_0 + o_P(1)$ with $\Sigma_0 > 0$, and*

$$\psi_n^* - \psi_0 = (P_n - P_{Q_0, g_n}) \left(d_{Y, \beta_0}^* + D_W(Q_{\beta_0}^*) \right) + o_P(1/\sqrt{n}). \quad (2.21)$$

Moreover, $\sqrt{\Sigma_n/n}(\psi_n^* - \psi_0)$ converges in law to the standard normal distribution.

Equality (2.21) is an asymptotic linear expansion of ψ_n^* under our targeted, adaptive sampling scheme. It is the key to the central limit theorem for $\sqrt{n}(\psi_n^* - \psi_0)$. It is important to note that this expansion does not rely on consistent estimation of $Q_{Y, 0}$, as we shall see in the derivations.

2.5 Example: targeted LASSO-based CARA RCT

In Sections 2.2, 2.3 and 2.4, we have presented a general framework for constructing and analyzing CARA RCTs using data-adaptive loss-based estimators for the nuisance parameters, coupled with the TMLE methodology to estimate the study parameter of interest. As described in Section 2.1, high-dimensional settings are increasingly common in clinical trials working with heterogeneous populations. A popular device in high-dimensional statistics, due to its computational feasibility and amenability to theoretical study, is the LASSO methodology. In a nutshell, the LASSO is a shrinkage and selection method for generalized regression models that optimizes a loss function of the regression coefficients subject to constraint on the L^1 norm. It was introduced by Tibshirani (1996) for obtaining estimators with fewer nonzero parameter values, thus effectively reducing the number of variables upon which the given solution is dependent. In this section, we illustrate the application of the proposed framework using the LASSO to estimate the conditional response and the optimal randomization scheme. The methodology introduced in Chambaz and van der Laan (2013) is a special case of this targeted LASSO-based CARA RCT.

For simplicity, we assume that all components of W are continuous. With a little extra work, discrete components could be handled as A is handled in (2.23).

Let $\ell^1 \equiv \{\beta \in \mathbb{R}^{\mathbb{N}} : \sum_{j \in \mathbb{N}} |\beta^j| < \infty\}$. Consider $\{b_n\}_{n \geq 1}$, $\{b'_n\}_{n \geq 1}$, $\{d_n\}_{n \geq 1}$, and $\{d'_n\}_{n \geq 1}$ four non-decreasing, possibly unbounded sequences over \mathbb{R}_+ and, for some $M, M' > 0$ and every $n \geq 1$, introduce the sets

$$B_n \equiv \{\beta \in \ell^1 : \|\beta\|_1 \leq \min(b_n, M) \text{ and } \forall j > d_n, \beta^j = 0\} \quad (2.22)$$

and B'_n defined like B_n with b'_n , d'_n and M' substituted for b_n , d_n , and M , respectively. Let $\{\phi_j : j \in \mathbb{N}\}$ be a uniformly bounded set of functions from \mathcal{W} to \mathbb{R} . Without loss of generality, we may assume that $\|\phi_j\|_\infty = 1$ for all $j \in \mathbb{N}$. By choice, the functions ϕ_j ($j \in \mathbb{N}$) share a common bounded support \mathcal{W} , and all belong to the class of sufficiently smooth functions, in the sense that there exists $\alpha > \dim(\mathcal{W})/2$ such that all partial derivatives up to order α of all ϕ_j exist and are uniformly bounded (see van der Vaart, 1998, Example 19.9).

For each β and $\omega \in \ell^1$, we denote $Q_{Y,\beta} : \mathcal{A} \times \mathcal{W} \rightarrow \mathbb{R}$ and $\gamma_\omega : \mathcal{A} \times \mathcal{W} \rightarrow \mathbb{R}$ the functions characterized by

$$Q_{Y,\beta}(A, W) \equiv \text{expit} \left(\sum_{j \in \mathbb{N}} (\beta^{2j} A + \beta^{2j+1} (1-A)) \phi_j(W) \right), \quad (2.23)$$

$$\gamma_\omega(1|W) = 1 - \gamma_\omega(0|W) \equiv \text{expit} \left(\sum_{j \in \mathbb{N}} \omega^j \phi_j(W) \right).$$

The LASSO-based CARA RCT design corresponds to a special choice of working models $\{\mathcal{Q}_{1,n}\}_{n \geq 1}$, $\{\mathcal{G}_{1,n}\}_{n \geq 1}$, and loss function L for $Q_{Y,0}$. We take $\mathcal{Q}_{1,n} \equiv \{Q_{Y,\beta} : \beta \in B_n\}$ with M a deterministic upper-bound on $|\text{logit}(Y)|$ and the quasi negative-log-likelihood loss function $L = L^{\text{kl}}$ (2.4). Note that the elements of $\mathcal{Q}_1 \equiv \cup_{n \geq 1} \mathcal{Q}_{1,n}$ are uniformly bounded away from 0 and 1. We also take $\mathcal{G}_{1,n} \equiv \{\gamma_\omega : \omega \in B'_n\}$. The elements of $\mathcal{G}_1 \equiv \cup_{n \geq 1} \mathcal{G}_{1,n}$ are randomization schemes uniformly bounded away from 0 and 1 by $\text{expit}(-M')$ and $\text{expit}(M')$, respectively ($M' \simeq 4.6$ provides the lower- and upper-bounds 0.01 and 0.99).

Based on \mathbf{O}_n , we estimate $Q_{Y,0}$ with $Q_{Y,\beta_{n+1}}$, where β_{n+1} is given in (2.8) (set $i = n + 1$ in the formula). Then we target g_0 with g_{n+1} given in (2.7) (set $i = n + 1$ in the formula), *i.e.*

$$g_{n+1} \in \arg \min_{\omega \in B'_n} \frac{1}{n} \sum_{i=1}^n \frac{L_{Q_{Y,\beta_n}}(\gamma_\omega)(O_i)}{g_i(A_i|W_i)}. \quad (2.24)$$

The minimization (2.8) with the constraint $\|\beta\|_1 \leq \min(b_n, M)$, see (2.22), can be rewritten as a minimization free of the latter constraint by adding a term of the form $\lambda_n \|\beta\|_1$ to the empirical criterion, where λ_n depends on b_n . Note that when d_n or d'_n is held constant and M or M' is infinite by choice, then (2.8) or (2.24) should be interpreted as a standard parametric procedure rather than as a LASSO.

Theorem 2.1 has the following corollary.

Corollary 2.2 (asymptotic study of the targeted LASSO-based CARA RCT). *Assume that **A1**, **A2**, **A3**, and **A4** are met. Then, the targeted LASSO-based CARA design converges in the sense that $\|g_n - g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ in probability as $n \rightarrow \infty$. Moreover, the TMLE ψ_n^* consistently estimates ψ_0 , and $\sqrt{n/\Sigma_n}(\psi_n^* - \psi_0)$ is approximately standard normally distributed, where Σ_n is the explicit estimator given in (2.20).*

The conditions on the $L1$ norms and dimensions (b_n, b'_n, d_n, d'_n) are in the assumptions on existence of the limits of the estimators. Beyond that, we only require that the basis functions be smooth. This corollary teaches us with minimal conditions on the smoothness of the basis functions, the targeted LASSO-based CARA RCT produces a convergent design and a consistent and asymptotically Gaussian estimator for the study parameter.

2.6 Simulation study

In this section, we exemplify the theoretical results from the previous sections with a brief simulation study. Specifically, we wish to (i) illustrate the robustness of the proposed TMLE estimator for the study parameter ψ_0 , under possibly grossly mis-specified conditional response models, (ii) show the use of data-adaptive LASSO estimators to learn the conditional response in the construction and analysis of the targeted CARA RCT, and (iii) evaluate the performances of the different strategies. The simulation study is conducted using R (R Core Team, 2014).

Data-generating distribution

Under Q_0 , $W = (U, V, Z_1, \dots, Z_{20})$ consists of 22 independent random variables, where U, Z_1, \dots, Z_{20} are all uniformly distributed on $[0, 1]$, and $V \in \{1, 2, 3\}$ is such that $V = 1$, $V = 2$ and $V = 3$ with probabilities $1/2$, $1/3$, and $1/6$, respectively. Moreover, under Q_0 and conditionally on (A, W) , Y is drawn from the Gamma distribution with conditional mean

$$Q_{Y,0}(A, W) \equiv 2AV + (1 - A)V/2$$

and conditional variance

$$\sigma_0^2(A, W) \equiv \left(\frac{AV}{3(1 + Z_1)} + \frac{4(1 - A)}{3(1 + Z_1)} \right)^2.$$

It is easy to check that $\psi_0 = 2.5$ and that the optimal randomization scheme g_0 is given by $g_0(1|W) \equiv V/(4 + V)$.

Loss functions and working models

To simplify the language, we refer to a model that accounts for the relevant covariates as a correctly specified model, even though the functional form may not be correct.

Estimation of the conditional response

Because Y is continuous and unbounded, we perform a linear transformation before the estimation procedures to scale Y within $(0, 1)$, then apply the reverse transformation to the final TMLE estimate of ψ_0 and the corresponding variance estimates. We use the quasi negative-log-likelihood loss function L^{kl} given by (2.4).

At sample size n , we consider two working models $\mathcal{Q}_{1,n}$ for the conditional response. One is the following mis-specified logistic regression model:

$$\mathcal{Q}_{1,n}^p \equiv \left\{ Q_{Y,\beta}^p(A, W) \equiv \text{expit}(\beta_1 A + \beta_2 U) : \beta \in \mathbb{R}^2 \right\}.$$

Contrary to what the notation suggests, it does not change as the sample size grows. It is fitted using the `glm` function in R with the weights as given in (2.8). Note that the model fails to take into account the covariate V which drives the response in the underlying data-generating process. The second one, denoted $\mathcal{Q}_{1,n}^\ell$, is a LASSO logistic working model. Let $d_n \equiv \min(20, \lfloor \sqrt{n}/4 \rfloor)$. If n is such that $d_n \leq 5$, then $\mathcal{Q}_{1,n}^\ell$ consists of

$$Q_{Y,\beta}^\ell(A, W) \equiv \text{expit} \left(\beta(A, U, Z_1, \dots, Z_{d_n}, AU, AZ_1, \dots, AZ_{d_n})^\top \right) \quad (\text{all } \beta \in B_n \equiv \mathbb{R}^{2d_n+3}).$$

If n is such that $d_n > 5$, then $\mathcal{Q}_{1,n}^\ell$ consists of

$$Q_{Y,\beta}^\ell(A, W) \equiv \text{expit} \left(\beta(A, U, V, Z_1, \dots, Z_{d_n}, AU, AV, AZ_1, \dots, AZ_{d_n})^\top \right) \quad (\text{all } \beta \in B_n \equiv \mathbb{R}^{2d_n+5}).$$

The resulting sequence of working models is non-decreasing in sample size. The models is fitted using the `cv.glmnet` function from the package `glmnet` (Friedman, Hastie, and Tibshirani, 2010), with weights given in (2.8) and the option `"lambda.1se"`.

Estimation of the optimal randomization scheme g_0

We also consider two working models $\mathcal{G}_{1,n} = \mathcal{G}_1$ for the optimal randomization scheme. The first one, denoted \mathcal{G}_1^m , is a mis-specified logistic model given by

$$g_\beta^m(A = 1 | W) \equiv \text{expit}(\beta_0 + \beta_1 U) \quad (\text{all } \beta \in \mathbb{R}^2).$$

The second one, denoted \mathcal{G}_1^c , is a correctly specified logistic model given by

$$g_\beta^c(A = 1 | W) \equiv \text{expit}(\beta_0 + \beta_1 U + \beta_2 V) \quad (\text{all } \beta \in \mathbb{R}^3).$$

The models are fitted using numerical methods to optimize the user-chosen adaptation criterion in (2.6). We implement this fitting using the `optim` function with a quasi-Newton method (`method="BFGS"`). To satisfy the boundedness conditions, the resulting probability estimates are pre-specified to be truncated to lie within $[0.05, 0.95]$. However, in the actual simulation runs, all estimates lying comfortably within this interval, and hence no truncation took place.

Study designs

For each pair of working models for the conditional response and for the optimal randomization scheme, we construct a CARA RCT by initializing at a sample size of $n = 300$, and then sequentially recruiting patients in blocks of size 200, up to $n = 3100$. For the initial sample of $n = 300$, treatment is randomly assigned based on the balanced randomization scheme g^b . Subsequently, given n observations, we estimate the conditional response and use this to construct the treatment randomization scheme g_{n+1} used for the next block of 200 patients. We also use this conditional response estimate and the sequence of randomization schemes used so far to obtain a TMLE estimate ψ_n^* of ψ_0 .

In addition to these CARA RCTs, we also consider a fixed design RCT with treatment randomly assigned based on the balanced randomization scheme g^b . We obtain the corresponding TMLE estimates by fluctuating the initial conditional response estimates based on the logistic model $\{Q_{Y,\beta}^p : \beta \in \mathbb{R}^2\}$.

Results

For each trial design proposed in Section 2.6, we run 500 independent simulated trials. Three figures summarize the results of the simulation study. Each of them consists of two similar graphics, the LHS graphic corresponding to the simulated trials based on the mis-specified model \mathcal{G}_1^m for the optimal randomization scheme, and the RHS graphic to the simulated trials based on the correctly specified model \mathcal{G}_1^c . The subtitles “A~U” and “A~U+V” are the R formulas that encode for \mathcal{G}_1^m and \mathcal{G}_1^c , respectively.

Figure 2.1 depicts the performance of ψ_n^* in terms of bias (first row), sample variance (second row) and mean squared error (MSE, third row). We note that, despite the mis-specified response models, all TMLE estimators are consistent for the treatment effect parameter ψ_0 . It appears that the LASSO-based estimator may converge at a faster rate. This may be due to its increased efficiency (*i.e.*, smaller sample variance) and more aggressive bias reduction. Recall that the optimality criterion for our adaptive randomization aims at maximizing efficiency of the trial through the minimization of the asymptotic variance of the estimators. The increased efficiency of the LASSO-based CARA RCT, despite a larger working model for the conditional response (increasing with sample size),

suggests that a flexible data-adaptive response model coupled with CARA design could indeed better achieve the optimality criterion, compared to a CARA design based on a parametric response model, at least in situations where the parametric model fails to account for important confounding variables. We also note that, under the data-generating process described in Section 2.6, the working model for the optimal randomization scheme has little effect on the efficiency of the TMLE estimators. Yet, comparing the LHS and RHS graphics in Figure 2.1 suggests that \mathcal{G}_1^m , the smaller, mis-specified model for the optimal randomization scheme allows for slightly more aggressive bias reduction at smaller sample sizes than \mathcal{G}_1^c , its larger, correctly specified counterpart.

Let us turn now to the coverage of our CLT-based, 95%-confidence intervals (CIs). The empirical coverage probabilities are depicted in Figure 2.2. On the one hand, we see that the empirical coverages are often below the nominal coverage when using the mis-specified working model \mathcal{G}_1^m for the optimal randomization scheme and either $\mathcal{Q}_{1,n}^p$ or $\mathcal{Q}_{1,n}^\ell$ as working models for the conditional response (LHS graphic in Figure 2.2). On the other hand, the coverage improves drastically when using the correctly specified working model \mathcal{G}_1^c for the optimal randomization scheme and either $\mathcal{Q}_{1,n}^p$ or $\mathcal{Q}_{1,n}^\ell$ as working models for the conditional response (RHS graphic in Figure 2.2). For a more precise assessment, we frame the coverage evaluation in terms of hypotheses testing. For a given design (and its resulting CLT-based CIs) and at each intermediate sample size n , let C be the number of times in the 500 simulations when the CI covers the parameter of interest ψ_0 . The random variable C is distributed from the Binomial distribution with parameter $(500, \pi)$. For a given significance level $0 < \alpha < 1$, introduce the null hypotheses $H_0^{1-\alpha} : \pi \geq 1 - \alpha$ and its one-sided alternative $H_1^{1-\alpha} : \pi < 1 - \alpha$. We perform one-sided tests of $H_0^{1-\alpha}$ against $H_1^{1-\alpha}$ and display the p -values for $\alpha = 5\%$ (Figure 2.3, first row) and $\alpha = 6\%$ (Figure 2.3, second row). On the one hand, the LHS graphic in Figure 2.3 reveals that 95% coverage is often not guaranteed when using the mis-specified working model $\mathcal{G}_{1,n}^m$ for the optimal randomization scheme and either $\mathcal{Q}_{1,n}^p$ or $\mathcal{Q}_{1,n}^\ell$ as working models for the conditional response, but also that 94% coverage cannot be ruled out. On the other hand, the RHS graphic in Figure 2.3 suggests that 95% coverage cannot be ruled out when using the correctly specified model $\mathcal{G}_{1,n}^c$ for the optimal randomization scheme and either $\mathcal{Q}_{1,n}^p$ or $\mathcal{Q}_{1,n}^\ell$ as working models for the conditional response.

2.7 Discussion

We have presented in this article a new group-sequential CARA RCT design and corresponding analytical procedure that admits the use of flexible data-adaptive techniques. The proposed method extends the work of Chambaz and van der Laan (2013) by provid-

ing robust inference of the study parameter under the most general settings. Our framework adopts a loss-based approach in estimating the optimal randomization scheme, and hence can target general optimality criteria that may not have a closed-form solutions. Moreover, our use of loss-based data-adaptive estimation over general classes of functions (which may change with sample size), both in constructing the treatment randomization schemes and in predicting the unknown conditional response, may potentially improve the randomization adaptation towards the optimality criterion.

The covariate adjustments take place in the estimation of the response, which has an effect in both constructing the randomization schemes, each depends on covariates through these outcome estimators, and in the estimation of the effect parameter itself, which directly uses the response estimator. Therefore adjustments can contribute to efficiency of the estimator in two folds: accounting for modifiers of the treatment effect during treatment assignment, as well as moving towards a consistent estimator of the response and hence achieving the efficiency bound.

We established that, under appropriate entropy conditions on the classes of functions, the resulting sequence of randomization schemes converges to a fixed scheme, and the proposed treatment effect estimator is consistent (even under a mis-specified response model), asymptotically Gaussian, giving rise to valid confidence intervals of given asymptotic levels. Moreover, the limiting randomization scheme coincides with the unknown optimal randomization scheme when, simultaneously, the response model is correctly specified and the optimal randomization scheme belongs to the limit of the user-supplied classes of randomization schemes. We illustrated the applicability of these general theoretical results with a LASSO-based CARA RCT. In this example, both the response model and the optimal treatment randomization are estimated using a sequence of LASSO logistic models that may increase with sample size. It follows immediately from our general theorems that this LASSO-based CARA RCT converges to a fixed randomization scheme and yields consistent and asymptotically Gaussian effect estimates, under minimal conditions on the smoothness of the basis functions in the LASSO logistic models.

We conducted a simulation study to evaluate the performance of the proposed methods. It confirmed the robustness of the TMLE estimators under mis-specified response models. Coverage of the CLT-based confidence intervals are assessed through by hypotheses testing. Overall there is no evidence (across 500 independent simulations) that the 95%-confidence intervals would have coverages that are less than 94%. In addition, we do observe improved coverage when using the correct working model for the optimal randomization scheme. In this work, we have provided an empirical assessment of the Type 1 error rate, but that is not often done in simulation studies on estimators of novel experiment designs. As preserving Type 1 error rate is a strong emphasis for many regulatory bodies, a more thorough simulation study should be performed to assess the empirical coverage of the proposed, as well as other traditional designs, under various experimen-

tal conditions. In this simulation study, the increased efficiency of CARA design with a LASSO-based response model, compared to the CARA (or balanced) design with a parametric response model, demonstrates that the use of data-adaptive response models can indeed more effectively steer the adaptation towards the optimality criterion (which was chosen to be efficiency in our example). More comprehensive empirical studies are needed to generalize these facts to other simulation scenarios.

We will soon make available a R package to allow interested readers to test the procedure. In the future, we will also consider alternative strategies to randomly assign successive patients to the treatment arms in such a way that the overall empirical conditional distribution of treatment given baseline covariates be as close as possible to the current best estimator of the targeted optimal randomization scheme. This will require both new theoretical developments and simulation studies.

2.8 Appendix

The expression “ $a \lesssim b$ ” means that there exists a universal, positive constant c such that $a \leq c \times b$. We use $\mathbf{1}\{\mathcal{C}\}$ to denote the indicator function of the set \mathcal{C} . We denote the uniform norm of a real-valued operator Π on \mathcal{F} as $\|\Pi\|_{\mathcal{F}} \equiv \sup_{f \in \mathcal{F}} |\Pi(f)|$. Given two measurable functions f, λ of (O, Z) and the random variable $\Lambda = \lambda(O, Z)$, we find it convenient to use shorthand notation $P_{Q_0, g} f \Lambda \equiv E_{P_{Q_0, g}}(f(O, Z)\Lambda)$ and $P_n f \Lambda \equiv E_{P_n}(f(O, Z)\Lambda) = n^{-1} \sum_{i=1}^n f(O_i, Z_i)\lambda(O_i, Z_i)$. From here onward, the uncountable supremum is interpreted as the essential supremum.

Section 2.8 presents the proofs of Propositions 2.1, 2.2, 2.3, 2.4, Corollary 2.1, Theorem 2.1 and Corollary 2.2. Technical results underpinning the proofs of Section 2.8 are gathered in Section 2.8.

Main proofs

Proof of Proposition 2.1. We apply Lemma 2.4 with $\Theta \equiv \mathcal{Q}_1$, $\Theta_n \equiv \mathcal{Q}_{1, n}$, d the distance induced on Θ by the norm $\|\cdot\|_{2, P_{Q_0, g^r}}$, \mathbf{M} and \mathbf{M}_n characterized over Θ by $\mathbf{M}(Q_Y) \equiv P_{Q_0, g^r} L(Q_Y)$ and $\mathbf{M}_n(Q_Y) \equiv P_n L(Q_Y) g^r / Z = n^{-1} \sum_{i=1}^n L(Q_Y)(O_i) g^r(A_i | W_i) / Z_i$. Assumption **A2** implies that **(a)** and **(b)** from Lemma 2.4 are met. It remains to prove that **(c)** also holds or, in other terms, that $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{Q}_{1, n}} = o_P(1)$.

For any $Q_Y \in \Theta$, characterize $\ell(Q_Y)$ by setting $\ell(Q_Y)(O, Z) \equiv L(Q_Y)(O) g^r(A|W) / Z$. Then we can rewrite $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{Q}_{1, n}}$ as follows:

$$\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{Q}_{1, n}} = \|P_n \ell - P_{Q_0, g^r} L\|_{\mathcal{Q}_{1, n}} = \|(P_n - P_{Q_0, \mathbf{g}_n}) \ell\|_{\mathcal{Q}_{1, n}} = \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\ell(\mathcal{Q}_{1, n})}.$$

The dominated ratio property implies that $J(1, \ell(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, g^r}}) = O(J(1, L(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, g^r}})) = o(\sqrt{n})$, by **A5**. Since $\ell(\mathcal{Q}_1)$ is uniformly bounded by construction, Lemma 2.8 applies and yields $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\ell(\mathcal{Q}_{1,n})} = o_P(1)$.

Thus, we can apply Lemma 2.4. It yields that $\|Q_{Y, \beta_n} - Q_{Y, \beta_0}\|_{2, P_{Q_0, g^r}} = o_P(1)$, which is the desired result. \square

The next proof goes along similar lines.

Proof of Proposition 2.2. We apply Lemma 2.4 with $\Theta \equiv \mathcal{G}_1$, $\Theta_n \equiv \mathcal{G}_{1,n}$, d the distance induced on Θ by the norm $\|\cdot\|_{2, Q_{W,0}}$. Over Θ , we define $\mathbf{M}(g) \equiv P_{Q_0, g^r} L_{Q_{Y, \beta_0}}(g)/g^r$ and $\mathbf{M}_n(g) \equiv P_n L_{Q_{Y, \beta_n}}(g)/Z = n^{-1} \sum_{i=1}^n L_{Q_{Y, \beta_n}}(g)(O_i)/Z_i$. Assumption **A3** implies that **(a)** and **(b)** from Lemma 2.4 are met. It remains to prove that **(c)** also holds or, in other terms, that $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{G}_{1,n}} = o_P(1)$.

Let ℓ and ℓ_n be characterized over \mathcal{G}_1 by $\ell(g)(O, Z) \equiv L_{Q_{Y, \beta_0}}(g)(O)/Z$ on the one hand and $\ell_n(g)(O, Z) \equiv L_{Q_{Y, \beta_n}}(g)(O)/Z$ on the other hand. A simple decomposition and the triangle inequality yield the following inequality:

$$\begin{aligned} \|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{G}_{1,n}} &= \|(P_n \ell - P_{Q_0, g^r} L_{Q_{Y, \beta_0}}/g^r) + P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}} \\ &\leq \|P_n \ell - P_{Q_0, g^r} L_{Q_{Y, \beta_0}}/g^r\|_{\mathcal{G}_{1,n}} + \|P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}} \\ &= \|(P_n - P_{Q_0, \mathbf{g}_n})\ell\|_{\mathcal{G}_{1,n}} + \|P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}} \\ &= \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\ell(\mathcal{G}_{1,n})} + \|P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}}. \end{aligned} \quad (2.25)$$

Consider the first RHS term in (2.25). Because Y and Q_{Y, β_0} are bounded, and because \mathcal{G}_1 is bounded away from 0 and 1 by construction, it holds that $J(1, \ell(\mathcal{G}_{1,n}), \|\cdot\|_{2, Q_{W,0}}) = O(J(1, 1/\mathcal{G}_{1,n}, \|\cdot\|_{2, Q_{W,0}})) = o(\sqrt{n})$ by **A6**. Since $\ell(\mathcal{G}_1)$ is uniformly bounded, Lemma 2.8 applies and yields $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\ell(\mathcal{G}_{1,n})} = o_P(1)$.

We now turn to the second RHS term in (2.25). Note $|L^{\text{ls}}(Q_{Y, \beta_n}) - L^{\text{ls}}(Q_{Y, \beta_0})| \lesssim |Q_{Y, \beta_n} - Q_{Y, \beta_0}|$ because Y is bounded and \mathcal{G}_1 is uniformly bounded. This justifies the second inequality below, the first one being a consequence of the uniform boundedness of $\mathcal{G}_{1,n}$, and the last one a consequence of the fact that g^r is bounded away from 0:

$$\begin{aligned} \|P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}} &\lesssim P_n |L^{\text{ls}}(Q_{Y, \beta_n}) - L^{\text{ls}}(Q_{Y, \beta_0})|/Z \\ &\lesssim P_n |Q_{Y, \beta_n} - Q_{Y, \beta_0}|/Z \\ &= P_{Q_0, \mathbf{g}_n} |Q_{Y, \beta_n} - Q_{Y, \beta_0}|/Z + (P_n - P_{Q_0, \mathbf{g}_n}) |Q_{Y, \beta_n} - Q_{Y, \beta_0}|/Z \\ &\lesssim P_{Q_0, g^r} |Q_{Y, \beta_n} - Q_{Y, \beta_0}| + (P_n - P_{Q_0, \mathbf{g}_n}) |Q_{Y, \beta_n} - Q_{Y, \beta_0}|/Z. \end{aligned}$$

Cauchy-Schwarz inequality implies that $P_{Q_0, g^r} |Q_{Y, \beta_n} - Q_{Y, \beta_0}| \leq \|Q_{Y, \beta_n} - Q_{Y, \beta_0}\|_{2, P_{Q_0, g^r}} = o_P(1)$ by Proposition 2.1, whose assumptions are met. For any $Q_Y \in \mathcal{Q}_1$, introduce $h(Q_Y)$

characterized by $h(Q_Y)(O, Z) \equiv |Q_{Y, \beta_n}(A, W) - Q_{Y, \beta_0}(A, W)|/Z$. Obviously,

$$|(P_n - P_{Q_0, \mathbf{g}_n})|Q_{Y, \beta_n} - Q_{Y, \beta_0}|/Z| \leq \|(P_n - P_{Q_0, \mathbf{g}_n})h\|_{\mathcal{Q}_{1,n}} = \|P_n - P_{Q_0, \mathbf{g}_n}\|_{h(\mathcal{Q}_{1,n})}.$$

Since \mathcal{Q}_1 and \mathcal{G}_1 are uniformly bounded away from 0 and 1 by construction, it holds that $h(\mathcal{Q}_1)$ is uniformly bounded and that $J(1, h(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, \mathbf{g}^r}}) = O(J(1, \{Q_Y - Q_{Y, \beta_0} : Q_Y \in \mathcal{Q}_{1,n}\}, \|\cdot\|_{2, P_{Q_0, \mathbf{g}^r}}) = o(\sqrt{n})$ by **A5**. Therefore, Lemma 2.8 applies and yields $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{h(\mathcal{Q}_{1,n})} = o_P(1)$.

We thus have showed that both $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\ell(\mathcal{G}_{1,n})} = o_P(1)$ and $\|P_n(\ell_n - \ell)\|_{\mathcal{G}_{1,n}} = o_P(1)$, hence $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{G}_{1,n}} = o_P(1)$ in light of (2.25). Consequently, we can apply Lemma 2.4. It yields that $\|g_n - g_0^*\|_{2, Q_{W,0}} = o_P(1)$, which is the desired result. \square

Proof of Corollary 2.1. Since \mathcal{G}_1 is uniformly bounded, $\|g_n - g_0^*\|_{2, Q_{W,0}} = o_P(1)$ implies $\|g_n - g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ in L^1 as $n \rightarrow \infty$. Since (i) $1/g_n - 1/g_0^* = (g_0^* - g_n)/g_n g_0^*$, and (ii) \mathcal{G}_1 is uniformly bounded away from 0 and 1, $\|1/g_n - 1/g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ follows from $\|g_n - g_0^*\|_{2, Q_{W,0}} \rightarrow 0$, both in probability and in L^1 as $n \rightarrow \infty$. Consider now the L^1 -convergence of $\|n^{-1} \sum_{i=1}^n g_i - g_0^*\|_{2, Q_{W,0}}$. By convexity,

$$E \left(\left\| \frac{1}{n} \sum_{i=1}^n g_i - g_0^* \right\|_{2, Q_{W,0}} \right) \leq \frac{1}{n} \sum_{i=1}^n E (\|g_i - g_0^*\|_{2, Q_{W,0}}).$$

We already know that $E(\|g_n - g_0^*\|_{2, Q_{W,0}}) = o(1)$; Cesaro's lemma yields $n^{-1} \sum_{i=1}^n E(\|g_i - g_0^*\|_{2, Q_{W,0}}) = o(1)$. From this, we deduce that $\|n^{-1} \sum_{i=1}^n g_i - g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ in L^1 as $n \rightarrow \infty$. This implies that the convergence also holds in probability because \mathcal{G}_1 is uniformly bounded. Likewise,

$$E \left(\left\| \frac{1}{n} \sum_{i=1}^n 1/g_i - 1/g_0^* \right\|_{2, Q_{W,0}} \right) \leq \frac{1}{n} \sum_{i=1}^n E (\|1/g_i - 1/g_0^*\|_{2, Q_{W,0}}),$$

where $E(\|1/g_n - 1/g_0^*\|_{2, Q_{W,0}}) = o(1)$ is already known. Thus, the same argument as above yields that $\|n^{-1} \sum_{i=1}^n 1/g_i - 1/g_0^*\|_{2, Q_{W,0}} \rightarrow 0$ in L^1 and in probability as $n \rightarrow \infty$. This completes the proof. \square

Proof of Proposition 2.3. This is a three-part proof. First, we show that $|\varepsilon_n - \varepsilon_0| = o_P(1)$. Second, we prove that $\|Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*\|_{2, P_{Q_0, \mathbf{g}^r}} = o_P(1)$. Third, we show that $\Psi(Q_{\beta_0}^*) = \psi_0$, then that ψ_n^* consistently estimates ψ_0 .

We apply (van der Vaart, 1998, Theorem 5.9) (substituting \mathbf{M}_n and \mathbf{M} for Ψ_n and Ψ) with $\Theta \equiv \mathcal{E}$, d the Euclidean distance, \mathbf{M} and \mathbf{M}_n characterized over Θ by $\mathbf{M}(\varepsilon) =$

$P_{Q_0, g_0^*} D_Y(Q_{Y, \beta_0}(\varepsilon), g_0^*)$, and $\mathbf{M}_n(\varepsilon) = P_n D_Y(Q_{Y, \beta_n}(\varepsilon), g_n) g_n / Z$, see (2.14) and (2.9) for the definitions of $Q_{Y, \beta_0}(\varepsilon)$ and $Q_{Y, \beta_n}(\varepsilon)$. From the differentiability of $\varepsilon \mapsto L^{\text{kl}}(Q_{Y, \beta}(\varepsilon))$, validity of the differentiation under the integral sign, and definition of ε_0 (2.15), we deduce that $\mathbf{M}(\varepsilon_0) = 0$. By definition of ε_n (2.10), $\mathbf{M}_n(\varepsilon_n) = 0$ too. Assumption **A4** implies that the second condition of the theorem is met. Therefore it suffices to check that the first one holds too, *i.e.* to prove that $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{E}} = o_P(1)$.

Introduce $\mathcal{F} = \{f_\varepsilon : \varepsilon \in \mathcal{E}\}$ with $f_\varepsilon(O, Z) \equiv (2A - 1)(Y - Q_{Y, \beta_0}(\varepsilon)(A, W))/Z$ for each $\varepsilon \in \mathcal{E}$. We start with the following derivation:

$$\begin{aligned} \|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{E}} &= \sup_{\varepsilon \in \mathcal{E}} \left| P_n \left(f_\varepsilon + \frac{2A - 1}{Z} (Q_{Y, \beta_0}(\varepsilon) - Q_{Y, \beta_n}(\varepsilon)) \right) - P_{Q_0, \mathbf{g}_n} f_\varepsilon \right| \\ &\leq \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}} + \sup_{\varepsilon \in \mathcal{E}} \left| P_n \frac{2A - 1}{Z} (Q_{Y, \beta_0}(\varepsilon) - Q_{Y, \beta_n}(\varepsilon)) \right|. \end{aligned} \quad (2.26)$$

Consider the first RHS term in (2.26). Set $\varepsilon_1, \varepsilon_2 \in \mathcal{E}$. Because the expit function is 1-Lipschitz and \mathcal{G}_1 is uniformly bounded, it holds that $\|f_{\varepsilon_1} - f_{\varepsilon_2}\|_\infty \lesssim |\varepsilon_1 - \varepsilon_2|$. Since \mathcal{E} is a bounded set by construction, the uniformly bounded, parametric class \mathcal{F} satisfies $J(1, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}}) < \infty$ (see van der Vaart, 1998, Example 19.7). Consequently, we can apply Lemma 2.8 (with a fixed class) and conclude that $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}} = o_P(1)$.

The second term in the RHS of (2.26) is upper-bounded by $\Delta_n \equiv \sup_{\varepsilon \in \mathcal{E}} P_n |Q_{Y, \beta_0}(\varepsilon) - Q_{Y, \beta_n}(\varepsilon)| / Z$. Since (i) expit is 1-Lipschitz, (ii) $\mathcal{Q}_{1, n}$ is bounded away from 0 and 1, and logit is Lipschitz on any compact subset of $]0, 1[$, it holds that

$$\begin{aligned} \Delta_n &\leq \sup_{\varepsilon \in \mathcal{E}} P_n \left| \text{logit}(Q_{Y, \beta_0}) - \text{logit}(Q_{Y, \beta_n}) + \varepsilon(H(g_n) - H(g_0^*)) \right| / Z \\ &\lesssim P_n |Q_{Y, \beta_0} - Q_{Y, \beta_n}| / Z + P_n |1/g_n - 1/g_0^*| / Z \\ &= (P_n - P_{Q_0, \mathbf{g}_n}) |Q_{Y, \beta_0} - Q_{Y, \beta_n}| / Z + P_{Q_0, \mathbf{g}_n} |Q_{Y, \beta_0} - Q_{Y, \beta_n}| / Z \\ &\quad + (P_n - P_{Q_0, \mathbf{g}_n}) |1/g_n - 1/g_0^*| / Z + P_{Q_0, \mathbf{g}_n} |1/g_n - 1/g_0^*| / Z. \end{aligned} \quad (2.27)$$

While studying the second RHS term of (2.25) in the proof of Proposition 2.2, we proved the following facts: $P_{Q_0, \mathbf{g}_n} |Q_{Y, \beta_0} - Q_{Y, \beta_n}| / Z \lesssim P_{Q_0, g^r} |Q_{Y, \beta_0} - Q_{Y, \beta_n}| = o_P(1)$ and $(P_n - P_{Q_0, \mathbf{g}_n}) |Q_{Y, \beta_0} - Q_{Y, \beta_n}| / Z = o_P(1)$ (the assumptions of Proposition 2.2 are met here too). Therefore, it only remains to study the two rightmost terms in the RHS of (2.27). Since \mathcal{G}_1 is uniformly bounded away from 0, $(P_n - P_{Q_0, \mathbf{g}_n}) |1/g_n - 1/g_0^*| / Z = O(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{1/\mathcal{G}_{1, n}})$. Moreover, Lemma 2.8 applies because $1/\mathcal{G}_1$ is uniformly bounded and **A5** is met, hence $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{1/\mathcal{G}_{1, n}} = o_P(1)$ and $(P_n - P_{Q_0, \mathbf{g}_n}) |1/g_n - 1/g_0^*| / Z = o_P(1)$. Finally,

$$P_{Q_0, \mathbf{g}_n} |1/g_n - 1/g_0^*| / Z \lesssim P_{Q_0, g^r} |1/g_n - 1/g_0^*| \leq \|1/g_n - 1/g_0^*\|_{2, Q_{W, 0}} = o_P(1)$$

by Cauchy-Schwarz and Corollary 2.1, whose assumptions are met here too. In summary, $\Delta_n = o_P(1)$.

We have show that the RHS expression in (2.26) converges to 0 in probability as $n \rightarrow \infty$, hence $\|\mathbf{M}_n - \mathbf{M}\|_{\mathcal{E}} = o_P(1)$. Thus, all assumptions of Lemma 2.4 hold, from which we deduce that ε_n converges to ε_0 as $n \rightarrow \infty$. This completes the first part of the proof.

Let $\mathcal{Q}_Y \times \mathcal{G} \times \mathcal{E}$ be equipped with the norm $\|(Q_Y, g, \varepsilon)\| = \|Q_Y\|_{2, P_{Q_0, g^r}} + \|g\|_{2, Q_{W,0}} + |\varepsilon|$. Propositions 2.1, 2.2 and the first part of the proof imply that $(Q_{Y, \beta_n}, g_n, \varepsilon_n)$ converges to $(Q_{Y, \beta_0}, g_0^*, \varepsilon_0)$ in probability wrt $\|\cdot\|$ as $n \rightarrow \infty$. Let $f : \mathcal{Q}_Y \times \mathcal{G} \times \mathcal{E} \rightarrow \mathcal{Q}_Y$ be characterized by

$$f(Q_Y, g, \varepsilon)(O) \equiv \text{expit}(\text{logit}(Q_Y(A, W)) + \varepsilon(2A - 1)/g(A|W)) \quad (2.28)$$

Set $(Q_{Y,1}, g_1, \varepsilon_1), (Q_{Y,2}, g_2, \varepsilon_2) \in \mathcal{Q}_Y \times \mathcal{G} \times \mathcal{E}$. Because (i) expit is 1-Lipschitz, (ii) $\mathcal{Q}_{1,n}$ is bounded away from 0 and 1, and logit is Lipschitz on any compact subset of $]0, 1[$, (iii) \mathcal{G}_1 is uniformly bounded away from 0, (iv) \mathcal{E} is a bounded set, it holds that

$$\begin{aligned} & \|f(Q_{Y,1}, g_1, \varepsilon_1) - f(Q_{Y,2}, g_2, \varepsilon_2)\|_{2, P_{Q_0, g^r}} \\ & \leq \|\text{logit}(Q_{Y,1}) - \text{logit}(Q_{Y,2})\|_{2, P_{Q_0, g^r}} + \|\varepsilon_2(1/g_1 - 1/g_2)\|_{2, Q_{W,0}} + \|(\varepsilon_1 - \varepsilon_2)/g_1\|_{2, Q_{W,0}} \\ & \lesssim \|Q_{Y,1} - Q_{Y,2}\|_{2, P_{Q_0, g^r}} + \|g_1 - g_2\|_{2, Q_{W,0}} + |\varepsilon_1 - \varepsilon_2| = \|(Q_{Y,1}, g_1, \varepsilon_1) - (Q_{Y,2}, g_2, \varepsilon_2)\| \end{aligned}$$

(f is Lipschitz). Therefore, the convergence $\|(Q_{Y, \beta_n}, g_n, \varepsilon_n) - (Q_{Y, \beta_0}, g_0^*, \varepsilon_0)\| = o_P(1)$ and equalities $Q_{Y, \beta_n}^* = f(Q_{Y, \beta_n}, g_n, \varepsilon_n)$, $Q_{Y, \beta_0}^* = f(Q_{Y, \beta_0}, g_0^*, \varepsilon_0)$, entail $\|Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} = o_P(1)$, hence our first claim. This completes the second part of the proof.

The second claim follows from the double-robustness of D . Indeed, it follows from the first part of this proof that $\mathbf{M}(\varepsilon_0) = P_{Q_0, g_0^*} D_Y(Q_{Y, \beta_0}^*, g_0^*) = 0$, and $P_{Q_0, g_0^*} D_W(Q_{\beta_0}^*) = 0$ from the definitions of Ψ and D_W , hence $P_{Q_0, g_0^*} D(P_{Q_{\beta_0}^*, g_0^*}) = 0$. Thus, Lemma 2.1 guarantees that $\Psi(Q_{\beta_0}^*) = \Psi(Q_0)$ since P_{Q_0, g_0^*} and $P_{Q_{\beta_0}^*, g_0^*}$ share the same g_0^* . We now turn to the third and last claim. For both $\beta = \beta_0$ and $\beta = \beta_n$, introduce $q_{Y, \beta}^*$ characterized by

$$q_{Y, \beta}^*(W) \equiv Q_{Y, \beta}^*(1, W) - Q_{Y, \beta}^*(0, W). \quad (2.29)$$

Define also $Q_{\beta_n}^{\sim} \equiv (Q_{W,0}, Q_{Y, \beta_n}^*)$ and $\psi_n^{\sim} \equiv \Psi(Q_{\beta_n}^{\sim})$. Since g^r is bounded away from 0, the Cauchy-Schwarz inequality yields

$$\begin{aligned} |\psi_n^{\sim} - \psi_0| &= |\psi_n^{\sim} - \Psi(Q_{\beta_0}^*)| = |P_{Q_0, g^r}(Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*)(2A - 1)/g^r| \\ &\lesssim P_{Q_0, g^r} |Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*| \leq \|Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} = o_P(1). \end{aligned}$$

Furthermore, $\psi_n^* - \psi_n^{\sim} = (P_n - P_{Q_0, g_n})(q_{Y, \beta_n}^* - q_{Y, \beta_0}^*) + (P_n - P_{Q_0, g_n})q_{Y, \beta_0}^*$. By similar arguments as before, we establish that $(P_n - P_{Q_0, g_n})(q_{Y, \beta_n}^* - q_{Y, \beta_0}^*) = o_P(1)$. In addition, the law

of large numbers (for independent, identically distributed random variables, since q_{Y,β_0}^* is a bounded function of W only) guarantees that $(P_n - P_{Q_0, \mathbf{g}_n})q_{Y,\beta_0}^* = o_P(1)$. In summary, $\psi_n^* - \psi_0 = (\psi_n^* - \psi_n^\sim) + (\psi_n^\sim - \psi_0) = o_P(1)$, as stated. This completes the proof. \square

The asymptotic linear expansion (2.21) in Proposition 2.4 is a by-product of the exact linear expansion that we state and prove below. Recall the definitions of $d_{Y,\beta}^*$ and $q_{Y,\beta}^*$ ($\beta = \beta_0$ or $\beta = \beta_n$) given in (2.18) and (2.29).

Lemma 2.2 (exact linear expansion of ψ_n^*). *It follows from the definition of ψ_n^* that*

$$\psi_n^* - \psi_0 = -P_{Q_0, g_0^*} D(P_{Q_{\beta_n}^*, g_0^*}) \quad (2.30)$$

$$\begin{aligned} &= (P_n - P_{Q_0, \mathbf{g}_n})(d_{Y, \beta_0}^* + D_W(Q_{\beta_0}^*)) \\ &\quad + (P_n - P_{Q_0, \mathbf{g}_n}) \left((d_{Y, \beta_n}^* - d_{Y, \beta_0}^*) + (q_{Y, \beta_n}^* - q_{Y, \beta_0}^*) \right). \end{aligned} \quad (2.31)$$

Proof of Lemma 2.2. Consider (2.30). By Lemma 2.1, D decomposes as $D(P_{Q_{\beta_n}^*, g_0^*}) = D_Y(Q_{Y, \beta_n}^*, g_0^*) + D_W(Q_{\beta_n}^*)$. Define $q_{Y,0}(W) \equiv Q_{Y,0}(1, W) - Q_{Y,0}(0, W)$. Firstly, we note $P_{Q_0, g_0^*} D_W(Q_{\beta_n}^*) = P_{Q_0, g_0^*} q_{Y, \beta_n}^* - \psi_n^*$. Secondly, $P_{Q_0, g_0^*} D_Y(Q_{Y, \beta_n}^*, g_0^*) = P_{Q_0, g_0^*} (2A - 1)(Y - Q_{Y, \beta_n}^*)/g_0^* = P_{Q_0, g_0^*} (q_{Y,0} - q_{Y, \beta_n}^*)$. Adding these two equalities yields $P_{Q_0, g_0^*} D(P_{Q_{\beta_n}^*, g_0^*}) = P_{Q_0, g_0^*} q_{Y,0} - \psi_n^* = \psi_0 - \psi_n^*$, which is the desired result.

We now turn to (2.31). Denote P_{n, \mathbf{g}_n} the empirical distribution of \mathbf{O}_n weighted by $g_n(A_i | W_i) / g_i(A_i | W_i)$. By construction of the fluctuation (2.9) and definition of ε_n (2.10), it holds that $P_{n, \mathbf{g}_n} D_Y(Q_{Y, \beta_n}^*, g_n) = 0$. Moreover, (2.11) can be rewritten as $P_n D_W(Q_{\beta_n}^*) = 0$. Therefore, (2.30) is equivalent to

$$\psi_n^* - \psi_0 = (P_n - P_{Q_0, g_0^*}) D_W(Q_{\beta_n}^*) + \left(P_{n, \mathbf{g}_n} D_Y(Q_{Y, \beta_n}^*, g_n) - P_{Q_0, g_0^*} D_Y(Q_{Y, \beta_n}^*, g_0^*) \right). \quad (2.32)$$

Adding and subtracting $(P_n - P_{Q_0, g_0^*}) D_W(Q_{\beta_0}^*)$ to the first term in the RHS expression of (2.32) implies

$$\begin{aligned} (P_n - P_{Q_0, g_0^*}) D_W(Q_{\beta_n}^*) &= (P_n - P_{Q_0, g_0^*}) D_W(Q_{\beta_0}^*) + (P_n - P_{Q_0, g_0^*}) (D_W(Q_{\beta_n}^*) - D_W(Q_{\beta_0}^*)) \\ &= (P_n - P_{Q_0, g_0^*}) D_W(Q_{\beta_0}^*) + (P_n - P_{Q_0, g_0^*}) (q_{Y, \beta_n}^* - q_{Y, \beta_0}^*) \\ &= (P_n - P_{Q_0, g_n}) D_W(Q_{\beta_0}^*) + (P_n - P_{Q_0, g_n}) (q_{Y, \beta_n}^* - q_{Y, \beta_0}^*), \end{aligned} \quad (2.33)$$

where the last equality is valid because $D_W(Q_{\beta_0}^*)$, q_{Y, β_n}^* , q_{Y, β_0}^* are functions of W only. As

for the second term in the RHS expression of (2.32), it equals

$$\begin{aligned}
& \frac{1}{n} \sum_{i=1}^n \left(\frac{g_n(A_i|W_i)}{g_i(A_i|W_i)} \frac{2A_i - 1}{g_n(A_i|W_i)} (Y_i - Q_{Y,\beta_n}^*(A_i, W_i)) - P_{Q_0, g_0^*} \frac{2A - 1}{g_0^*(A|W)} (Y - Q_{Y,\beta_n}^*) \right) \\
&= \frac{1}{n} \sum_{i=1}^n \left(\frac{2A_i - 1}{g_i(A_i|W_i)} (Y_i - Q_{Y,\beta_n}^*(A_i, W_i)) - P_{Q_0, g_i} \frac{2A - 1}{g_i(A|W)} (Y - Q_{Y,\beta_n}^*) \right) \\
&= (P_n - P_{Q_0, \mathbf{g}_n}) d_{Y,\beta_n}^* \\
&= (P_n - P_{Q_0, \mathbf{g}_n}) d_{Y,\beta_0}^* + (P_n - P_{Q_0, \mathbf{g}_n}) (d_{Y,\beta_n}^* - d_{Y,\beta_0}^*). \tag{2.34}
\end{aligned}$$

The equalities (2.32), (2.33) and (2.34) imply (2.31). \square

It appears that the second term in the RHS expression of (2.31) is asymptotically negligible at rate \sqrt{n} . Indeed,

Lemma 2.3. *It holds that $(P_n - P_{Q_0, \mathbf{g}_n}) \left((d_{Y,\beta_n}^* - d_{Y,\beta_0}^*) + (q_{Y,\beta_n}^* - q_{Y,\beta_0}^*) \right) = o_P(1/\sqrt{n})$.*

Proof of Lemma 2.3. The key to this proof is Lemma 2.10.

Introduce $\mathcal{Q}_{1,n}^* \equiv \{f(Q_{Y,\beta}, g, \varepsilon) : Q_{Y,\beta} \in \mathcal{Q}_{1,n}, g \in \mathcal{G}_{1,n}, \varepsilon \in \mathcal{E}\}$, where f is given by (2.28), and set $\delta > 0$. The elements of $\mathcal{Q}_{1,n}^*$ are uniformly bounded away from 0 and 1. By **A5**, **A6** and Lemma 2.11, the bracketing numbers $N(\delta, \text{logit}(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, g^r}})$ and $N(\delta, 1/\mathcal{Q}_{1,n}, \|\cdot\|_{2, P_{Q_0, g^r}})$ are finite. Obviously, the bracketing number $N(\delta, \mathcal{E}, |\cdot|)$ is finite too. Choose arbitrarily three collections of δ -brackets of smallest possible cardinality that cover $\text{logit}(\mathcal{Q}_{1,n})$, $1/\mathcal{G}_{1,n}$, and \mathcal{E} . Given $f(Q_{Y,\beta}, g, \varepsilon) \in \mathcal{Q}_{1,n}^*$, let $[l_Q, u_Q]$, $[l_g, u_g]$ and $[l_\varepsilon, u_\varepsilon]$ be δ -brackets from these collections and containing $\text{logit}(Q_{Y,\beta})$, $1/g$ and ε , respectively. We can assume without loss of generality that the uniform lower- and upper-bounds of $\text{logit}(\mathcal{Q}_{1,n})$ (respectively, $1/\mathcal{G}_{1,n}$) are also lower- and upper-bounds on l_Q , u_Q , (respectively, l_g , u_g). We can also assume that $|l_\varepsilon|, |u_\varepsilon| \leq \sup_{\varepsilon \in \mathcal{E}} |\varepsilon|$. Characterize λ and γ by setting $\lambda(O) \equiv A l_\varepsilon H(u_g)(O) + (1-A) u_\varepsilon H(l_g)(O)$ and, similarly, $\gamma(O) \equiv A u_\varepsilon H(l_g)(O) + (1-A) l_\varepsilon H(u_g)(O)$. Then $[\text{expit}(l_Q + \lambda), \text{expit}(u_Q + \gamma)]$ is a bracket containing $f(Q_{Y,\beta}, g, \varepsilon)$. Since expit is 1-Lipschitz, it follows that

$$(\text{expit}(u_Q + \gamma) - \text{expit}(l_Q + \lambda))^2 \leq ((u_Q - l_Q) + (\gamma - \lambda))^2 \leq 2(u_Q - l_Q)^2 + 2(\gamma - \lambda)^2$$

where $(\gamma - \lambda)^2 \lesssim (u_\varepsilon - l_\varepsilon)^2 + (H(u_g) - H(l_g))^2 \lesssim (u_\varepsilon - l_\varepsilon)^2 + (u_g - l_g)^2$. Consequently, there exists a universal constant $c \geq 1$ such that $[\text{expit}(l_Q + \lambda), \text{expit}(u_Q + \gamma)]$ be a $c\delta$ -bracket. Thus, $N(\delta, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\delta/c, \text{logit}(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{Q_0, g^r}}) \times N(\delta/c, 1/\mathcal{G}_{1,n}, \|\cdot\|_{2, P_{Q_0, g^r}}) \times N(\delta/c, \mathcal{E}, |\cdot|)$ hence, by Lemma 2.11, $J(\delta, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}}) \lesssim J(\delta, \mathcal{Q}_{1,n}, \|\cdot\|_{2, P_{Q_0, g^r}}) + J(\delta, \mathcal{G}_{1,n}, \|\cdot\|_{2, P_{Q_0, g^r}}) + J(\delta, \mathcal{E}, |\cdot|)$. Therefore, **A5*** and **A6*** imply that if

$\delta_n = o(1)$ then $J(\delta_n, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$ as well. Now, we use this to prove the lemma.

For each $Q_Y \in \mathcal{Q}_{1,n}^*$, characterize $d_Y(Q_Y)$ by setting $d_Y(Q_Y)(O, Z) \equiv (2A - 1)(Y - Q_Y(A, W))/Z$. By uniform boundedness of $\cup_{n \geq 1} \mathcal{Q}_{1,n}^*$, Y and Z , the existence of a sequence of envelope functions satisfying (a) in Lemma 2.10 is granted. Moreover, Lemma 2.11 yields that there exists $c > 0$ such that $J(\delta, d_Y(\mathcal{Q}_{1,n}^*), \|\cdot\|_{2, P_{Q_0, g^r}}) \leq cJ(\delta, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}})$ for all $\delta > 0$. Thus, $\delta_n = o(1)$ implies $J(\delta_n, d_Y(\mathcal{Q}_{1,n}^*), \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$, and condition (b) in Lemma 2.10 is met too. Now, the convergence $\|Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} = o_P(1)$, established in Proposition 2.3, implies $P_{Q_0, g^r}(d_Y(Q_{Y, \beta_n}^*) - d_Y(Q_{Y, \beta_0}^*))^2 = o_P(1)$ by Cauchy-Schwarz, since $|d_Y(Q_{Y, \beta_n}^*) - d_Y(Q_{Y, \beta_0}^*)| \lesssim |Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*|$. We apply Lemma 2.10 to obtain $\sqrt{n}(P_n - P_{Q_0, \mathbf{g}_n})(d_{Y, \beta_n}^* - d_{Y, \beta_0}^*) = o_P(1)$.

Now, for each $Q_Y \in \mathcal{Q}_{1,n}^*$, define $q_Y(Q_Y)(W) \equiv Q_Y(1, W) - Q_Y(0, W)$. Choose a collection of $N(\delta, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}})$ δ -brackets $[l_k, u_k]$ covering $\mathcal{Q}_{1,n}^*$. For a given $Q_Y \in \mathcal{Q}_{1,n}^*$, assume without loss of generality that $Q_Y \in [l_1, u_1]$ and define $l'_1(W) \equiv l_1(1, W) - u_1(0, W)$ and $u'_1(W) \equiv u_1(1, W) - l_1(0, W)$. It holds that $q_Y(Q_Y) \in [l'_1, u'_1]$ and $P_{Q_0, g^r}(u'_1 - l'_1)^2 \leq 2\delta^2/c$ for $0 < c \equiv \min(\inf g^r, 1 - \sup g^r) < 1$. Therefore, $N(\delta, q_Y(\mathcal{Q}_{1,n}^*), \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\sqrt{2/c}\delta, \mathcal{Q}_{1,n}^*, \|\cdot\|_{2, P_{Q_0, g^r}})$, and thus $J(\delta_n, q_Y(\mathcal{Q}_{1,n}^*), \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$ whenever $\delta_n = o(1)$. Therefore, condition (b) in Lemma 2.10 is met. Condition (a) in the same lemma is also met since $\cup_{n \geq 1} q_Y(\mathcal{Q}_{1,n}^*)$ is uniformly bounded. Moreover, $\|Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} = o_P(1)$ implies $\|q_Y(Q_{Y, \beta_n}^*) - q_Y(Q_{Y, \beta_0}^*)\|_{2, P_{Q_0, g^r}} = o_P(1)$ as $P_{Q_0, g^r}(q_Y(Q_{Y, \beta_n}^*) - q_Y(Q_{Y, \beta_0}^*))^2 \leq 2P_{Q_0, g^r}(Q_{Y, \beta_n}^*(1, W) - Q_{Y, \beta_0}^*(1, W))^2 + 2P_{Q_0, g^r}(Q_{Y, \beta_n}^*(1, W) - Q_{Y, \beta_0}^*(1, W))^2$ and, for both $a = 0, 1$, $P_{Q_0, g^r}(Q_{Y, \beta_n}^*(a, W) - Q_{Y, \beta_0}^*(a, W))^2 = P_{Q_0, g^r}(Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*)^2 \mathbf{1}\{A = a\}/g^r(a|W) \lesssim P_{Q_0, g^r}(Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*)^2$ because g^r is bounded away from 0 and 1. We apply Lemma 2.10 to obtain $\sqrt{n}(P_n - P_{Q_0, \mathbf{g}_n})(q_{Y, \beta_n}^* - q_{Y, \beta_0}^*) = o_P(1)$.

This completes the proof. \square

The proof of Proposition 2.4 is now at hand.

Proof of proposition 2.4. We first note that (2.21) follows straightforwardly from Lemmas 2.2 and 2.3.

Set $f_0 \equiv d_{Y, \beta_0}^* + D_W(Q_{\beta_0}^*)$ and $f_n \equiv d_{Y, \beta_n}^* + D_W(Q_{\beta_n}^*)$. With this notation, $\Sigma_0 = P_{Q_0, \mathbf{g}_0} f_0^2$, $\Sigma_n = P_n f_n^2$. Introduce also $S_n \equiv P_{Q_0, \mathbf{g}_n} f_0^2$. For either $(f, \beta) = (f_0, \beta_0)$ or $(f, \beta) = (f_n, \beta_n)$,

it holds that

$$\begin{aligned}
P_{Q_0, \mathbf{g}_n} f^2 &= \frac{1}{n} \sum_{i=1}^n P_{Q_0, g_i} f^2 \\
&= P_{Q_0, g_0^*} \left(D_W(Q_\beta^*)^2 + 2D_Y(Q_{Y, \beta}^*, g_0^*) D_W(Q_\beta^*) \right) + \frac{1}{n} \sum_{i=1}^n P_{Q_0, g_0^*} \frac{(Y - Q_{Y, \beta}^*)^2}{g_0^* g_i} \\
&= P_{Q_0, g_0^*} \left(D_W(Q_\beta^*)^2 + 2D_Y(Q_{Y, \beta}^*, g_0^*) D_W(Q_\beta^*) \right) + P_{Q_0, g_0^*} \frac{(Y - Q_{Y, \beta}^*)^2}{g_0^*} \frac{1}{n} \sum_{i=1}^n 1/g_i.
\end{aligned}$$

Now, because $(Y - Q_{Y, \beta}^*)^2 \leq 1$ and g_0^* is bounded away from 0 and 1, the Cauchy-Schwarz inequality yields

$$\begin{aligned}
|P_{Q_0, \mathbf{g}_n} f^2 - P_{Q_0, g_0^*} f^2| &= \left| P_{Q_0, g_0^*} \frac{(Y - Q_{Y, \beta}^*)^2}{g_0^*} \left(\frac{1}{n} \sum_{i=1}^n 1/g_i - 1/g_0^* \right) \right| \\
&\lesssim P_{Q_0, g_0^*} \left| \frac{1}{n} \sum_{i=1}^n 1/g_i - 1/g_0^* \right| \leq \left\| \frac{1}{n} \sum_{i=1}^n 1/g_i - 1/g_0^* \right\|_{2, Q_{W, 0}}. \quad (2.35)
\end{aligned}$$

Thus, taking $f = f_0$ and applying corollary 2.1, we obtain $E(S_n) = \Sigma_0 + o(1)$ and $S_n = \Sigma_0 + o_P(1)$ (Note that $\Sigma_0 > 0$ by **A1**. Let us show now that $\Sigma_n = \Sigma_0 + o_P(1)$ by proving $\Sigma_n - S_n = o_P(1)$). We use the following decomposition:

$$\begin{aligned}
\Sigma_n - S_n &= (P_n - P_{Q_0, \mathbf{g}_n})(f_n^2 - f_0^2) + (P_n - P_{Q_0, \mathbf{g}_n})f_0^2 + P_{Q_0, \mathbf{g}_n}(f_n^2 - f_0^2) \\
&= (P_n - P_{Q_0, \mathbf{g}_n})(f_n^2 - f_0^2) + (P_n - P_{Q_0, \mathbf{g}_n})f_0^2 + P_{Q_0, g_0^*}(f_n^2 - f_0^2) + o_P(1) \quad (2.36)
\end{aligned}$$

where the second equality holds because $P_{Q_0, \mathbf{g}_n} f^2 = P_{Q_0, g_0^*} f^2 + o_P(1)$ for both $f = f_0$ and $f = f_n$ (by (2.35) and Corollary 2.1). Because f_0 and all f_n 's ($n \geq 1$) are uniformly bounded, the first term in the RHS expression of (2.36) satisfies

$$\begin{aligned}
|(P_n - P_{Q_0, \mathbf{g}_n})(f_n^2 - f_0^2)| &\lesssim |(P_n - P_{Q_0, \mathbf{g}_n})(f_n - f_0)| \\
&= |(P_n - P_{Q_0, \mathbf{g}_n})(d_{Y, \beta_n}^* - d_{Y, \beta_0}^*) + (q_{Y, \beta_n}^* - q_{Y, \beta_0}^*)| = o_P(1/\sqrt{n})
\end{aligned}$$

by Lemma 2.3 (see (2.29) for the definition of $q_{Y, \beta}^*$). Since f_0 is bounded, the Kolmogorov strong law of large numbers (Sen and Singer, 1993, Theorem 2.4.2) guarantees that the second term in the RHS expression of (2.36) converges to 0 P -almost-surely, hence $(P_n - P_{Q_0, \mathbf{g}_n})f_0^2 = o_P(1)$. Consider now the third term in the RHS expression of (2.36). Note that $(f_n - f_0)(O, Z) = (2A - 1)(Q_{Y, \beta_0}^* - Q_{Y, \beta_n}^*)(A, W)/Z + (q_{Y, \beta_n}^* - q_{Y, \beta_0}^*)(W) - (\psi_n^* - \psi_0)$, hence $|f_n - f_0| \lesssim |Q_{Y, \beta_0}^* - Q_{Y, \beta_n}^*| + |q_{Y, \beta_n}^* - q_{Y, \beta_0}^*| + |\psi_n^* - \psi_0|$ because Z is bounded away

from 0 and 1. Using again (i) that f_0 and all f_n 's ($n \geq 1$) are uniformly bounded, and (ii) the Cauchy-Schwarz inequality and the dominated ratio property, we get

$$\begin{aligned} |P_{Q_0, g_0^*}(f_n^2 - f_0^2)| &\lesssim |P_{Q_0, g_0^*}(f_n - f_0)| \\ &\lesssim P_{Q_0, g_0^*}|Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*| + P_{Q_0, g_0^*}|q_{Y, \beta_n}^* - q_{Y, \beta_0}^*| + |\psi_n^* - \psi_0| \\ &\lesssim \|Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} + \|q_{Y, \beta_n}^* - q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} + |\psi_n^* - \psi_0|. \end{aligned}$$

We know that $\|Q_{Y, \beta_n}^* - Q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} = o_P(1)$ by Proposition 2.1, we showed at the end of the proof of Lemma 2.3 that this implies $\|q_{Y, \beta_n}^* - q_{Y, \beta_0}^*\|_{2, P_{Q_0, g^r}} = o_P(1)$, and Proposition 2.3 guarantees that $\psi_n^* - \psi_0 = o_P(1)$. Consequently, $|P_{Q_0, g_0^*}(f_n^2 - f_0^2)| = o_P(1)$. We have thus proven that all terms in the RHS expression of (2.36) are $o_P(1)$, hence $\Sigma_n - S_n = o_P(1)$ and $\Sigma_n = \Sigma_0 + o_P(1)$, as we claimed earlier.

We show now that (2.21), which we rewrite here $\psi_n^* - \psi_0 = (P_n - P_{Q_0, g_n})f_0 + o_P(1/\sqrt{n})$, implies that $\sqrt{n/\Sigma_0}(\psi_n^* - \psi_0)$ converges in law to the standard normal distribution. This is a consequence of (Sen and Singer, 1993, Theorem 3.3.7) because (i) $S_n/E(S_n) - 1 = o_P(1)$, and (ii) for each $\alpha > 0$, $E(P_n f_0^2 \mathbf{1}\{f_0^2 \geq \alpha^2 n E(S_n)\}) = o(E(S_n))$ trivially holds since f_0 is bounded and $E(S_n) = \Sigma_0 + o(1)$ with $\Sigma_0 > 0$. Then Slutsky's lemma and $\Sigma_n = \Sigma_0 + o_P(1)$ yield the convergence in law of $\sqrt{n/\Sigma_n}(\psi_n^* - \psi_0)$ to the same limiting distribution. This completes the proof. \square

The proof of Corollary 2.2 boils down to (i) showing that **A5**, **A5***, **A6**, **A6*** are met and (ii) applying Theorem 2.1.

Proof of Corollary 2.2. We show below that **A5** and **A5*** are met. A parallel argument can be used to show that **A6** and **A6*** hold too. Since **A1**–**A4** are satisfied by assumption, Theorem 2.1 thus applies and yields the stated result.

Fix $\delta > 0$, a sequence $\{\delta_n\}_{n \geq 1}$ of positive numbers such that $\delta_n = o(1)$, and $n \geq 1$. By construction, the functions ϕ_j ($j \in \mathbb{N}$) all belong to a class \mathcal{C} of smooth functions over the bounded support \mathcal{W} such that all partial derivatives up to order $\alpha > \dim(\mathcal{W})/2$ of all $f \in \mathcal{C}$ exist and are uniformly bounded by a constant $C > 0$. By (van der Vaart, 1998, Example 19.9), it holds that $\log N(\delta, \mathcal{C}, \|\cdot\|_{2, P_{Q_0, g^r}}) \lesssim \delta^{-V}$ for $V \equiv \dim(\mathcal{W})/\alpha < 2$.

Note that $\mathcal{F} \equiv \{\sum_{j \in \mathbb{N}} \beta_j \phi_j = \sum_{j=0}^{d_n} \beta_j \phi_j : \beta \in B_n\}$ is a subset of \mathcal{C} , provided that the constant C in the definition of \mathcal{C} is large enough (if not, it suffices to replace C with MC , with M the constant involved in (2.22)). We apply three times Lemma 2.11 to obtain that $J(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}}) \gtrsim J(\delta, \text{logit}(\mathcal{Q}_{1, n}), \|\cdot\|_{2, P_{Q_0, g^r}}) \gtrsim J(\delta, \mathcal{Q}_{1, n}, \|\cdot\|_{2, P_{Q_0, g^r}}) \gtrsim J(\delta, L^{\text{kl}}(\mathcal{Q}_{1, n}), \|\cdot\|_{2, P_{Q_0, g^r}})$: from left to right, the inequalities follow from (i) the third claim of Lemma 2.11 with h, h' given by $h(O) \equiv A$ and $h'(O) \equiv (1 - A)$, (ii) from the sixth

claim with $\phi \equiv \text{expit}$, which is increasing and 1-Lipschitz, and (iii) from the seventh claim with h given by $h(O) \equiv Y$. Therefore,

$$\begin{aligned} J(\delta_n, L^{\text{kl}}(\mathcal{Q}_{1,n}), \|\cdot\|_{2, P_{\mathcal{Q}_0, g^r}}) &\lesssim J(\delta_n, \mathcal{Q}_{1,n}, \|\cdot\|_{2, P_{\mathcal{Q}_0, g^r}}) \\ &\lesssim J(\delta_n, \mathcal{C}, \|\cdot\|_{2, P_{\mathcal{Q}_0, g^r}}) \lesssim \int_0^{\delta_n} \varepsilon^{-V/2} d\varepsilon = o(1), \end{aligned}$$

and **A5*** is fulfilled. Choosing $\delta_n = 1/\sqrt{n}$ yields that **A5** is also fulfilled. This completes the proof. \square

Useful technical results

Convergence of M -estimators.

The following lemma is a simple adaptation of (van der Vaart and Wellner, 1996a, Corollary 3.2.3).

Lemma 2.4 (convergence of M -estimators). *Let \mathbf{M}_n be a real-valued, stochastic processes indexed by a metric space (Θ, d) , and let $\mathbf{M} : \Theta \rightarrow \mathbb{R}$ be a real-valued, deterministic function over Θ . Consider a sequence of subsets $\Theta_n \subset \Theta$ and the following assumptions:*

- (a) *There exists $\theta_0 \in \Theta$ such that $\mathbf{M}(\theta_0) < \inf_{\theta \notin T} \mathbf{M}(\theta)$ for every open set $T \subset \Theta$ containing θ_0 .*
- (b) *For each $n \geq 1$, there exists $\theta_n^* \in \Theta_n$ such that $\mathbf{M}(\theta_n^*) = \inf_{\theta \in \Theta_n} \mathbf{M}(\theta)$. Moreover, $\mathbf{M}(\theta_n^*) - \mathbf{M}(\theta_0) = o(1)$.*
- (c) *It holds that $\|\mathbf{M}_n - \mathbf{M}\|_{\Theta_n} = o_P(1)$.*

Under the above three assumptions, if $\theta_n \in \Theta_n$ satisfies $\mathbf{M}_n(\theta_n) - \mathbf{M}_n(\theta_n^) \leq 0$ for all $n \geq 1$, then $d(\theta_n, \theta_0) = o_P(1)$.*

Proof of Lemma 2.4. Set $n \geq 1$. By (a), it holds that

$$\begin{aligned} 0 &\leq \mathbf{M}(\theta_n) - \mathbf{M}(\theta_0) \\ &= (\mathbf{M}(\theta_n) - \mathbf{M}_n(\theta_n)) + (\mathbf{M}_n(\theta_n) - \mathbf{M}_n(\theta_n^*)) + (\mathbf{M}_n(\theta_n^*) - \mathbf{M}(\theta_n^*)) + (\mathbf{M}(\theta_n^*) - \mathbf{M}(\theta_0)). \end{aligned}$$

The above first and third RHS terms are both upper-bounded by $\|\mathbf{M}_n - \mathbf{M}\|_{\Theta_n}$. The second RHS term is non-positive by definition of θ_n . The fourth RHS term is $o(1)$ by (b). Thus, it actually holds that $0 \leq \mathbf{M}(\theta_n) - \mathbf{M}(\theta_0) \leq 2\|\mathbf{M}_n - \mathbf{M}\|_{\Theta_n} + o(1) = o_P(1)$ by (c).

Set $\varepsilon > 0$. By (a), there exists $\delta > 0$ such that $d(\theta_n, \theta_0) \geq \varepsilon$ implies $\mathbf{M}(\theta_n) - \mathbf{M}(\theta_0) \geq \delta$. Since we have shown that $\mathbf{M}(\theta_n) - \mathbf{M}(\theta_0) = o_P(1)$, we can therefore conclude that $d(\theta_n, \theta_0) = o_P(1)$ too. \square

Maximal inequalities and convergence of empirical processes.

In this article, we repeatedly exploit uniform laws of large numbers. They are derived from maximal inequalities for martingales by van Handel (2011) that also played an important role in (Chambaz and van der Laan, 2011a,c). For completeness, we now state these results.

Let $\phi : \mathbb{R} \rightarrow \mathbb{R}_+$ be such that $\phi(x) = e^x - x - 1$. Let \mathcal{F} be a class of measurable functions, $n \geq 1$ be an integer, $K > 0$ and $\delta > 0$ be two positive constants. For each $f \in \mathcal{F}$, $n(P_n - P_{Q_0, \mathbf{g}_n})f = \sum_{i=1}^n (f(O_i, Z_i) - P_{Q_0, \mathbf{g}_i} f)$ is a discrete martingale sum.

Set $N = N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, \mathbf{g}^r}})$, the δ -bracketing number of \mathcal{F} wrt $\|\cdot\|_{2, P_{Q_0, \mathbf{g}^r}}$. Following van Handel (2011), we define a $(\delta, n, \mathcal{F}, K)$ -bracketing set as a collection $\{(\Lambda_i^j, \Gamma_i^j) : i \leq n\}_{j \leq N}$ of random variables such that (i) for each $f \in \mathcal{F}$, there exists $j \leq N$ satisfying $\Lambda_i^j \leq f(O_i, Z_i) \leq \Gamma_i^j$ for all $i \leq n$, and (ii) for all $j \leq N$, $2K^2 n^{-1} \sum_{i=1}^n P_{Q_0, \mathbf{g}_i} \phi(|\Lambda_i^j - \Gamma_i^j|/K) \leq \delta^2$. Let $\mathcal{N}(\delta, n, \mathcal{F}, K)$ denote the cardinality of the smallest $(\delta, n, \mathcal{F}, K)$ -bracketing set. Finally, define for each $f \in \mathcal{F}$ the random variable $R_{n, K}(f) = 2K^2 n^{-1} \sum_{i=1}^n P_{Q_0, \mathbf{g}_i} \phi(|f|/K)$.

Lemma 2.5 (Proposition A.2 by van Handel (2011)). *There exists an universal constant $C > 0$ such that, for all $R > 0$,*

$$P \left(\sup_{f \in \mathcal{F}} \mathbf{1}\{R_{n, K}(f) \leq R\} \max_{i \leq n} \frac{i}{n} (P_i - P_{Q_0, \mathbf{g}_i}) f \geq \alpha \right) \leq 2 \exp \left(- \frac{n\alpha^2}{C^2(c_1 + 1)R} \right),$$

for any $\alpha, c_0, c_1 > 0$ such that $c_0^2 \geq C^2(c_1 + 1)$ and

$$\frac{c_0}{\sqrt{n}} \int_0^{\sqrt{R}} \sqrt{\log \mathcal{N}(\varepsilon, n, \mathcal{F}, K)} d\varepsilon \leq \alpha \leq \frac{c_1 R}{K}.$$

Lemma 2.6 (Corollary A.8 by van Handel (2011)). *Suppose the class \mathcal{F} is finite. For all $R > 0$ and any event C ,*

$$E \left(\max_{f \in \mathcal{F}} \mathbf{1}\{nR_{n, K}(f) \leq R\} \max_{i \leq n} i (P_i - P_{Q_0, \mathbf{g}_i}) f \right) \leq \sqrt{2R \log \left(1 + \frac{|\mathcal{F}|}{P(C)} \right)} + 8K \log \left(1 + \frac{|\mathcal{F}|}{P(C)} \right).$$

If, in addition, $\max_{f \in \mathcal{F}} \|f\|_\infty \leq U$, then K can be replaced with $U/3$ in the second term of the above RHS expression.

Importantly, van Handel (2011)'s proofs of Lemmas 2.5 and 2.6 remain valid when the class \mathcal{F} is allowed to depend on n . To use lemmas 2.5 and 2.6, it is necessary to get a grip on $\mathcal{N}(\delta, n, \mathcal{F}, K)$ and the random variables $R_{n, K}(f)$, $f \in \mathcal{F}$. The next lemma is helpful in this regard.

Recall that, by the dominated ratio property of \mathcal{G}_1 , $\|g/g^r\|_\infty \leq \kappa$ for all $g \in \mathcal{G}_1$.

Lemma 2.7 (L^2 -norm version of lemma 7 by Chambaz and van der Laan (2011c)). *Assume that $U \equiv \sup_{f \in \mathcal{F}} \|f\|_\infty$ is finite. Then, for all $f \in \mathcal{F}$, $R_{n,4U}(f) \leq 4/3n \sum_{i=1}^n P_{Q_0, g_i} |f|^2$. Moreover, it holds that $\mathcal{N}(\sqrt{2\kappa}\delta, n, \mathcal{F}, 4U) \leq N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$.*

Proof of Lemma 2.7. Set $f \in \mathcal{F}$, $i \leq n$, and $m \geq 2$. It holds that

$$P_{Q_0, g_i} |f|^m \leq U^{m-2} P_{Q_0, g_i} |f|^2 \leq \frac{m!}{2} U^{m-2} P_{Q_0, g_i} |f|^2.$$

Therefore, for $K = 4U$,

$$\begin{aligned} 2K^2 P_{Q_0, g_i} \phi(|f|/K) &= 2(4U)^2 \sum_{m \geq 2} \frac{P_{Q_0, g_i} |f|^m}{m!(4U)^m} \\ &\leq 2(4U)^2 \sum_{m \geq 2} \frac{\frac{m!}{2} U^{m-2} P_{Q_0, g_i} |f|^2}{m!(4U)^m} = 16 \sum_{m \geq 2} \frac{P_{Q_0, g_i} |f|^2}{4^m} = 4P_{Q_0, g_i} |f|^2/3. \end{aligned}$$

The monotone convergence theorem guarantees the first equality. Summing up the above inequalities for $i = 1, \dots, n$ yields the first bound.

Let $(\ell^j, u^j)_{j \leq N}$ be a set of $N = N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$ δ -brackets covering \mathcal{F} wrt $\|\cdot\|_{2, P_{Q_0, g^r}}$. Let $\Lambda_i^j = \max(\ell^j(O_i, Z_i), -U)$ and $\Gamma_i^j = \min(u^j(O_i, Z_i), U)$ for all $i \leq n, j \leq N$. Set $f \in \mathcal{F}$ and $j \leq N$ such that $f \in [\ell^j, u^j]$. Then, for all $i \leq n$, (i) $\Lambda_i^j \leq f(O_i, Z_i) \leq \Gamma_i^j$, (ii) $-U \leq \Lambda_i^j \leq \Gamma_i^j \leq U$, and (iii) $\ell^j \leq \Lambda_i^j \leq \Gamma_i^j \leq u^j$. Thus, for all $m \geq 2$,

$$\begin{aligned} P_{Q_0, g_i} |\Lambda_i^j - \Gamma_i^j|^m &\leq (2U)^{m-2} P_{Q_0, g_i} |\Lambda_i^j - \Gamma_i^j|^2 \leq (2U)^{m-2} \kappa P_{Q_0, g^r} |\Lambda_i^j - \Gamma_i^j|^2 \\ &\leq (2U)^{m-2} \kappa P_{Q_0, g^r} |\ell^j - u^j|^2 \leq (2U)^{m-2} \kappa \delta^2 \leq \frac{m!}{2} (2U)^{m-2} \kappa \delta^2. \end{aligned}$$

Consequently, still using $K = 4U$, it holds that

$$2K^2 P_{Q_0, g_i} \phi(|\Lambda_i^j - \Gamma_i^j|/4U) = 2(4U)^2 \sum_{m \geq 2} \frac{P_{Q_0, g_i} |\Lambda_i^j - \Gamma_i^j|^m}{m!(4U)^m} \leq 32U^2 \sum_{m \geq 2} \frac{\frac{m!}{2} (2U)^{m-2} \kappa \delta^2}{m!(4U)^m} = 2\kappa \delta^2.$$

Again, the monotone convergence theorem validates the first equality. Summing up the above inequalities for $i = 1, \dots, n$ yields $2K^2 n^{-1} \sum_{i=1}^n P_{Q_0, g_i} \phi(|\Lambda_i^j - \Gamma_i^j|/K) \leq 2\kappa \delta^2$, hence $\{(\Lambda_i^j, \Gamma_i^j) : i \leq n\}_{j \leq N}$ is a $(\sqrt{2\kappa}\delta, n, \mathcal{F}, K)$ -bracketing set and $\mathcal{N}(\sqrt{2\kappa}\delta, n, \mathcal{F}, 4U) \leq N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$. This completes the proof. \square

By combining Lemmas 2.5 and 2.7, we now establish a uniform law of large numbers.

Lemma 2.8. *Let $\{\mathcal{F}_n\}_{n \geq 1}$ be a sequence of sets of measurable functions such that $U \equiv \sup_{f \in \mathcal{U}_{n \geq 1}} \|f\|_\infty$ be finite. If $J(\sqrt{2/3\kappa}U, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(\sqrt{n})$, then for all $\alpha > 0$ there exists $c > 0$ and $n_0 \geq 1$ such that, for every $n \geq n_0$,*

$$P \left(\sup_{f \in \mathcal{F}_n} (P_n - P_{Q_0, \mathbf{g}_n}) f \geq \alpha \right) \leq 2e^{-nc}.$$

Consequently, $\sup_{f \in \mathcal{F}_n} |(P_n - P_{Q_0, \mathbf{g}_n})f|$ converges to 0 P -almost surely.

Lemma 2.8 modifies Theorem 8 in Chambaz and van der Laan (2011c) to use an L^2 -metric and allow the classes of functions to change with n .

Proof of Lemma 2.8. Set $\alpha > 0$, and let $K = 4U$, $R = 4/3U^2$, $c_1 = \alpha K/R$, $c_0 = C\sqrt{c_1 + 1}$, where C is the universal constant from Lemma 2.5. Note that $\sqrt{R/2\kappa} = \sqrt{2/3\kappa}U$. By assumption, there exists $n_0 \geq 1$ such that, for all $n \geq n_0$, $J(\sqrt{R/2\kappa}, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq \sqrt{n}\alpha/c_0\sqrt{2\kappa}$.

Set $n \geq n_0$. By Lemma 2.7, $\mathcal{N}(\sqrt{2\kappa}\delta, n, \mathcal{F}_n, 4U) \leq N(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})$. Therefore,

$$\begin{aligned} \frac{c_0}{\sqrt{n}} \int_0^{\sqrt{R}} \sqrt{\log \mathcal{N}(\varepsilon, n, \mathcal{F}_n, 4U)} d\varepsilon &\leq \frac{\sqrt{2\kappa}c_0}{\sqrt{n}} \int_0^{\sqrt{R/2\kappa}} \sqrt{\log N(\varepsilon, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})} d\varepsilon \\ &= \frac{\sqrt{2\kappa}c_0}{\sqrt{n}} J(\sqrt{R/2\kappa}, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq \alpha = \frac{c_1 R}{K}. \end{aligned}$$

Lemma 2.5 applies and yields here

$$P \left(\sup_{f \in \mathcal{F}_n} (P_n - P_{Q_0, \mathbf{g}_n}) f \geq \alpha \right) \leq P \left(\sup_{f \in \mathcal{F}_n} \max_{i \leq n} \frac{i}{n} (P_i - P_{Q_0, \mathbf{g}_i}) f \geq \alpha \right) \leq 2e^{-nc},$$

with $c = \alpha^2/c_0^2 R$. This completes the proof. \square

Lemma 2.5 also allows us to adapt the maximal inequality of (van der Vaart, 1998, Lemma 19.34), valid under independent, identically distributed sampling, to our targeted, adaptive sampling. We state and prove this result in lemma 2.9. We introduce the function Log given by $\text{Log}(x) \equiv \max(1, \log(x))$ (all $x > 0$).

Lemma 2.9. *Let \mathcal{F} be a class of measurable, real-valued functions and $\delta > 0$ be such that $P_{Q_0, g^r} f^2 \leq \delta^2$ for every $f \in \mathcal{F}$. Let F be an envelope function of \mathcal{F} . Define $a(\varepsilon) = \varepsilon / \sqrt{\text{Log} N(\varepsilon, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})}$ for all $\varepsilon > 0$. For each $n \geq 1$, it holds that*

$$\begin{aligned} & \sqrt{n}E(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}}) \\ & \lesssim J(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}}) + \sqrt{n}E(P_{Q_0, \mathbf{g}_n}F\mathbf{1}\{F > \sqrt{na}(\delta)\}) \end{aligned} \quad (2.37)$$

$$\leq J(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}}) + \sqrt{n}\kappa P_{Q_0, g^r}F\mathbf{1}\{F > \sqrt{na}(\delta)\}. \quad (2.38)$$

Proof of Lemma 2.9. The proof parallels that of (van der Vaart, 1998, Lemma 19.34).

Preliminary. Inequality (2.38) follows readily from (2.37) because $F\mathbf{1}\{F > \sqrt{na}(\delta)\}$ is non-negative and \mathcal{G}_1 is endowed with the dominated ratio property. To understand the sum of two terms on the RHS of (2.37), first note that $E(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}})$ is upper-bounded by

$$E\left(\sup_{f \in \mathcal{F}} |(P_n - P_{Q_0, \mathbf{g}_n})f\mathbf{1}\{F \leq \sqrt{na}(\delta)\}|\right) + E\left(\sup_{f \in \mathcal{F}} |(P_n - P_{Q_0, \mathbf{g}_n})f\mathbf{1}\{F > \sqrt{na}(\delta)\}|\right). \quad (2.39)$$

Now, for every $f \in \mathcal{F}$,

$$|(P_n - P_{Q_0, \mathbf{g}_n})f\mathbf{1}\{F > \sqrt{na}(\delta)\}| \leq (P_n + P_{Q_0, \mathbf{g}_n})F\mathbf{1}\{F > \sqrt{na}(\delta)\},$$

hence, by the tower rule, the second term in (2.39) is smaller than $E((P_n + P_{Q_0, \mathbf{g}_n})F\mathbf{1}\{F > \sqrt{na}(\delta)\}) = 2E(P_{Q_0, \mathbf{g}_n}F\mathbf{1}\{F > \sqrt{na}(\delta)\})$. Thus, to prove (2.37), it remains to show that \sqrt{n} times the first term in (2.39) is smaller than $J(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, up to a universal, multiplicative constant.

Before proceeding, note that $N(\varepsilon, \{f\mathbf{1}\{F \leq \sqrt{na}(\delta)\} : f \in \mathcal{F}\}, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\varepsilon, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$ for all $\varepsilon > 0$. Therefore, we may assume that $\sup_{f \in \mathcal{F}} \|f\|_{\infty} \leq \sqrt{na}(\delta)$. What follows is based on a chaining technique to replace \mathcal{F} with a finite class.

Chaining. We now define a nested sequence of partitions on \mathcal{F} , then deduce a finite representation of \mathcal{F} from it. Fix q_0 such that $\delta \leq 2^{-q_0} \leq 2\delta$. For each integer $q \geq q_0$, denote $\tilde{N}_q \equiv N(2^{-q}, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$. Since $\varepsilon \mapsto N(\varepsilon, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$ is non-decreasing, it holds that

$$\sum_{q \geq q_0} 2^{-q} \sqrt{\text{Log } \tilde{N}_q} \lesssim \int_0^{\delta} \sqrt{\text{Log } N(\varepsilon, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})} d\varepsilon. \quad (2.40)$$

1. For each $q \geq q_0$, cover \mathcal{F} with \tilde{N}_q many brackets $[l_{q,i}, u_{q,i}]_{i \leq \tilde{N}_q}$ such that $P_{Q_0, g^r} \Delta_{q,i}^2 \leq 2^{-2q}$ for all $i \leq \tilde{N}_q$. Note that we may assume, without loss of generality, that $\Delta_{q,i} \equiv u_{q,i} - l_{q,i} \leq 2F \leq 2\sqrt{na}(\delta)$ for all $i \leq \tilde{N}_q$. Define $\mathcal{F}_{q,1} \equiv [l_{q,1}, u_{q,1}]$ then, recursively, $\mathcal{F}_{q,i} \equiv [l_{q,i}, u_{q,i}] \cap (\bigcup_{j < i} [l_{q,j}, u_{q,j}])^c$ for $2 \leq i \leq \tilde{N}_q$. We have our first partition: $\mathcal{F} = \bigcup_{i=1}^{\tilde{N}_q} \mathcal{F}_{q,i}$, which we call partition of \mathcal{F} at level q .

From the sequence of partitions $\{\{\mathcal{F}_{q,i} : i \leq \tilde{N}_q\}\}_{q \geq q_0}$, we derive a nested sequence of partitions as follows. The first partition is $\{\mathcal{F}_{q_0,i} : i \leq \tilde{N}_{q_0}\}$ itself. Then, recursively, at a level q such that $\{\mathcal{F}_{q,i} : i \leq \tilde{N}_q\}$ is not a successful refinement of $\{\mathcal{F}_{(q-1),i} : i \leq \tilde{N}_{q-1}\}$, we replace each partitioning set at level q by its intersection with all partitioning sets at level $(q-1)$. All partitioning sets derived in this fashion from $\mathcal{F}_{q,i}$ are associated with the same $\Delta_{q,i}$. For a given $q \geq q_0$, the possibly new partition consists of at most $N_q = \prod_{q'=q_0}^q \tilde{N}_{q'}$ partitioning sets. Using the inequality $\sqrt{\text{Log } N_q} \leq \sum_{q'=q_0}^q \sqrt{\text{Log } \tilde{N}_{q'}}$, we see that (2.40) is preserved in the sense that

$$\begin{aligned} \sum_{q \geq q_0} 2^{-q} \sqrt{\text{Log } N_q} &\leq \sum_{q \geq q_0} 2^{-q} \sum_{q'=q_0}^q \sqrt{\text{Log } \tilde{N}_{q'}} \\ &\lesssim \sum_{q \geq q_0} 2^{-q} \sqrt{\text{Log } \tilde{N}_q} \leq \int_0^\delta \sqrt{\text{Log } N(\varepsilon, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})} d\varepsilon. \end{aligned} \quad (2.41)$$

2. At each level $q \geq q_0$ and for each $\mathcal{F}_{q,i}$ ($i \leq N_q$), fix a representative $f_{q,i} \in \mathcal{F}_{q,i}$. For every $f \in \mathcal{F}$, if $f \in \mathcal{F}_{q,i}$, then we set $\pi_q f \equiv f_{q,i}$ and $\Delta_q f \equiv \Delta_{q,i}$. Introduce $a_{q_0} \equiv 2a(2^{-q_0}) = 2^{-q_0+1} / \sqrt{\text{Log } \tilde{N}_{q_0}}$ and, for each $q > q_0$, $f \in \mathcal{F}$, $B_q f \equiv 0$,

$$\begin{aligned} a_q &\equiv 2^{-q+1} / \sqrt{\text{Log } N_q}, \\ A_{q-1} f &\equiv \mathbf{1}\{\Delta_{q_0} f \leq \sqrt{na_{q_0}}, \dots, \Delta_{q-1} f \leq \sqrt{na_{q-1}}\}, \\ B_q f &\equiv \mathbf{1}\{\Delta_{q_0} f \leq \sqrt{na_{q_0}}, \dots, \Delta_{q-1} f \leq \sqrt{na_{q-1}}, \Delta_q f > \sqrt{na_q}\}. \end{aligned}$$

By nestedness of the sequence of partitions, $f \mapsto A_q f$ and $f \mapsto B_q f$ are constant over each $\mathcal{F}_{q,i}$ ($i \leq N_q$). Moreover, $B_q f + A_q f = A_{q-1} f$ for all $q > q_0$ and $f \in \mathcal{F}$. In addition, since $\varepsilon \mapsto a(\varepsilon)$ is non-decreasing, $\Delta_{q_0} f \leq 2F \leq 2\sqrt{na}(\delta) \leq 2\sqrt{na}(2^{-q_0}) = \sqrt{na_{q_0}}$, hence $A_{q_0} f = 1$.

Using these facts, any $f \in \mathcal{F}$ decomposes as

$$f = \pi_{q_0} f + \sum_{q \geq q_0+1} (f - \pi_q f) B_q f + \sum_{q \geq q_0+1} (\pi_q f - \pi_{q-1} f) A_{q-1} f. \quad (2.42)$$

To see this, note first that either (i) $B_q f = 0$ for all $q \geq q_0$, which implies, by recursion, that $A_q f = 1$ for all $q \geq q_0$, or (ii) there exists $q_1 \geq q_0$ such that $B_{q_1} f = 1$, in which case $B_q f = 0$ for all $q \geq q_0, q \neq q_1$, and $A_q f = 1$ for all $q_0 \leq q < q_1$, $A_q f = 0$ for all $q \geq q_1$. If (i) holds, then we deal with a telescopic sum and (2.42) boils down to $f = \pi_{q_0} f + \lim_{q \rightarrow \infty} \pi_q f - \pi_{q_0} f$. The above equality is valid because both $\pi_q f$ and f are in the bracket $[l_q, u_q]$, whose size $\|u_q - l_q\|_{2, P_{Q_0, g^r}} \rightarrow 0$ as $n \rightarrow \infty$. If (ii) holds, then $f = \pi_{q_0} f + (f - \pi_{q_1} f) + \sum_{q=q_0+1}^{q_1} (\pi_q f - \pi_{q-1} f)$ is evidently true.

Define $\mathcal{F}_a = \{\pi_{q_0}f/\sqrt{n} : f \in \mathcal{F}\}$, $\mathcal{F}_b = \{\sum_{q \geq q_0+1} (f - \pi_q f) B_q f / \sqrt{n} : f \in \mathcal{F}\}$, and $\mathcal{F}_c = \{\sum_{q \geq q_0+1} (\pi_q f - \pi_{q-1} f) A_{q-1} f / \sqrt{n} : f \in \mathcal{F}\}$. Each sum in the definition of \mathcal{F}_b consists of at most one single term. Each sum in the definition of \mathcal{F}_c is either finite, or telescopic, with a limit, in which case the dominated convergence theorem guarantees that $P_{Q_0, \mathbf{g}_n} \sum_{q \geq q_0+1} (\pi_q f - \pi_{q-1} f) A_{q-1} f = \sum_{q \geq q_0+1} P_{Q_0, \mathbf{g}_n} (\pi_q f - \pi_{q-1} f) A_{q-1} f$. Therefore, (2.42) yields

$$E(\|(P_n - P_{Q_0, \mathbf{g}_n})\|_{\mathcal{F}}) / \sqrt{n} \leq E(\|(P_n - P_{Q_0, \mathbf{g}_n})\|_{\mathcal{F}_a}) + E(\|(P_n - P_{Q_0, \mathbf{g}_n})\|_{\mathcal{F}_b}) + E(\|(P_n - P_{Q_0, \mathbf{g}_n})\|_{\mathcal{F}_c}). \quad (2.43)$$

We shall study in turn each term in the RHS expression of (2.43).

Class \mathcal{F}_a . For every $f \in \mathcal{F}$, (i) $|\pi_{q_0}f| \leq \sqrt{na}(\delta) \leq \sqrt{na}(2^{-q_0}) = \sqrt{na}a_{q_0}/2$, hence $\sup_{h \in \mathcal{F}_a} \|h\|_\infty \leq a_{q_0}/2$, and (ii) $P_{Q_0, \mathbf{g}^r}(\pi_{q_0}f)^2 \leq \delta^2$ (true by assumption). Apply Lemma 2.6 with $\mathcal{F} = \mathcal{F}_a$, C the whole probability space, $U = a_{q_0}/2$, $K = 4U$, $R = 4\kappa\delta^2/3$ (an upper-bound on $nR_{n, 4U}(\pi_{q_0}f/\sqrt{n})$ valid uniformly in $f \in \mathcal{F}$ by Lemma 2.7): it holds that

$$\begin{aligned} nE(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_a}) &\lesssim \delta \sqrt{\text{Log } N_{q_0}} + a_{q_0} \text{Log } N_{q_0} \\ &\leq 2^{-q_0} \sqrt{\text{Log } N_{q_0}} + 2^{-q_0+1} \frac{\text{Log } N_{q_0}}{\sqrt{\text{Log } N_{q_0}}} \\ &\leq \sum_{q \geq q_0} 2^{-q} \sqrt{\text{Log } N_q}. \end{aligned} \quad (2.44)$$

Class \mathcal{F}_b . For every $q > q_0$, $f \in \mathcal{F}$, $|f - \pi_q f| \leq \Delta_q f$ implies

$$|(P_n - P_{Q_0, \mathbf{g}_n})(f - \pi_q f)| \leq (P_n + P_{Q_0, \mathbf{g}_n})\Delta_q f \leq |(P_n - P_{Q_0, \mathbf{g}_n})\Delta_q f| + 2P_{Q_0, \mathbf{g}_n}\Delta_q f.$$

Thus, by using repeatedly the triangle inequality and the dominated convergence theorem, we obtain

$$\begin{aligned} &E(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_b}) \\ &\leq \sum_{q \geq q_0+1} E\left(\sup_{f \in \mathcal{F}} |(P_n - P_{Q_0, \mathbf{g}_n})\Delta_q f B_q f / \sqrt{n}|\right) + 2 \sum_{q \geq q_0+1} E\left(\sup_{f \in \mathcal{F}} P_{Q_0, \mathbf{g}_n} \Delta_q f B_q f / \sqrt{n}\right). \end{aligned} \quad (2.45)$$

Consider the first term in the RHS expression of (2.45). Fix $q > q_0$. Note that $f, f' \in \mathcal{F}_{q,i}$ implies $\Delta_q f B_q f = \Delta_q f' B_q f'$. So, the supremum $\sup_{f \in \mathcal{F}} |(P_n - P_{Q_0, \mathbf{g}_n})\Delta_q f B_q f / \sqrt{n}|$ is actually a maximum over a set of cardinality N_q . Moreover, for each $f \in \mathcal{F}$, (i) $0 \leq \Delta_q f B_q f \leq \Delta_{q-1} f B_q f \leq \sqrt{na}a_{q-1}$, hence $\sup_{h \in \mathcal{F}_b} \|h\|_\infty \leq a_{q-1}$, and (ii) $P_{Q_0, \mathbf{g}^r}(\Delta_q f B_q f)^2 \leq$

2^{-2q} . Apply Lemma 2.6 with $\mathcal{F} = \mathcal{F}_b$, C the whole probability space, $U = a_{q-1}$, $K = 4U$, $R = 4\kappa 2^{-2q}/3$ (an upper-bound on $nR_{n,4U}(\Delta_q f B_q f / \sqrt{n})$ valid uniformly in $f \in \mathcal{F}$ by Lemma 2.7): it holds that

$$\begin{aligned} nE \left(\sup_{f \in \mathcal{F}} |(P_n - P_{Q_0, \mathbf{g}_n}) \Delta_q f B_q f / \sqrt{n}| \right) &\lesssim 2^{-q} \sqrt{\text{Log } N_q} + a_{q-1} \text{Log } N_q \\ &= 2^{-q} \sqrt{\text{Log } N_q} + 2^{-q+2} \frac{\text{Log } N_q}{\sqrt{\text{Log } N_q}} \\ &\lesssim 2^{-q} \sqrt{\text{Log } N_q}. \end{aligned} \quad (2.46)$$

Consider now the second term in (2.45). Fix $q > q_0$ and $f \in \mathcal{F}$. Since $B_q f = 1$ only if $\sqrt{na_q} < \Delta_q f$, it follows that

$$\sqrt{na_q} P_{Q_0, \mathbf{g}_i} \Delta_q f B_q f \leq P_{Q_0, \mathbf{g}_i} (\Delta_q f)^2 B_q f \leq 2^{-2q}$$

for every $1 \leq i \leq n$. Therefore,

$$\sup_{f \in \mathcal{F}} P_{Q_0, \mathbf{g}_n} \Delta_q f B_q f / \sqrt{n} \leq 2^{-2q} / na_q \lesssim 2^{-q} \sqrt{\text{Log } N_q} / n. \quad (2.47)$$

By (2.45), summing up (2.46) and (2.47) over $q > q_0$ finally yields

$$nE(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_b}) \lesssim \sum_{q \geq q_0+1} 2^{-q} \sqrt{\text{Log } N_q}. \quad (2.48)$$

Class \mathcal{F}_c . Fix $q > q_0$. Note that $f, f' \in \mathcal{F}_{q,i}$ implies $(\pi_q f - \pi_{q-1} f) A_q f = (\pi_q f' - \pi_{q-1} f') A_q f'$. So, the supremum $\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_c}$ is actually a maximum over a set of cardinality N_q . Moreover, for each $f \in \mathcal{F}$, (i) $|\pi_q f - \pi_{q-1} f| A_{q-1} f \leq \Delta_{q-1} f A_{q-1} f \leq \sqrt{na_{q-1}}$, hence $\sup_{h \in \mathcal{F}_c} \|h\|_\infty \leq a_{q-1}$, from which we also deduce that (ii) $P_{Q_0, \mathbf{g}_r} ((\pi_q f - \pi_{q-1} f) A_{q-1} f)^2 \leq P_{Q_0, \mathbf{g}_r} (\Delta_q f)^2 \leq 2^{-2q}$. Therefore, the same reasoning as the one which lead us to (2.48) applies again, and we obtain

$$nE(\|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_c}) \lesssim \sum_{q \geq q_0+1} 2^{-q} \sqrt{\text{Log } N_q}. \quad (2.49)$$

Combining (2.43), (2.44), (2.48), (2.49), and (2.41) completes the proof. \square

To prove Proposition 2.4, we must study the convergence in probability of empirical processes indexed by estimated functions. Lemma 2.10 below provides sufficient conditions to derive such convergences. The version of this lemma under a i.i.d. sampling scheme is given by (van der Vaart and Wellner, 2007, Theorem 2.2). Here, we provide its extension to the current targeted adaptive sampling scheme. The proof of Lemma 2.10 hinges on Lemma 2.9.

Lemma 2.10 (convergence of empirical processes indexed by estimated functions). *For each $n \geq 1$, let $\mathcal{F}_n = \{f_{\theta, \eta} : \theta \in \Theta, \eta \in T_n\}$ be a class of measurable, real-valued functions, with envelope function F_n . Suppose the following holds:*

(a) *The sequence $\{F_n\}_{n \geq 1}$ satisfies the Lindeberg condition: $P_{Q_0, g^r} F_n^2 = O(1)$ and, for every $\delta > 0$, $P_{Q_0, g^r} F_n^2 \mathbf{1}\{F_n > \delta \sqrt{n}\} = o(1)$.*

(b) *If $\delta_n = o(1)$, then it holds that $J(\delta_n, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$.*

If $\eta_n \in T_n$ is such that $\sup_{\theta \in \Theta} P_{Q_0, g^r}(f_{\theta, \eta_n} - f_{\theta, \eta_0})^2 = o_P(1)$ for some $\eta_0 \in \cap_{p \geq 1} \cup_{n \geq p} T_n$, then $\sup_{\theta \in \Theta} |\sqrt{n}(P_n - P_{Q_0, \mathbf{g}_n})(f_{\theta, \eta_n} - f_{\theta, \eta_0})| = o_P(1)$.

Proof of lemma 2.10. Define the random class $\widetilde{\mathcal{F}}_n^0 \equiv \{f_{\theta, \eta_n} - f_{\theta, \eta_0} : \theta \in \Theta\}$. We wish to prove that $\sqrt{n} \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\widetilde{\mathcal{F}}_n^0} = o_P(1)$. Set arbitrarily $\alpha > 0, \varepsilon > 0$, and introduce $\mathcal{F}_n^0 \equiv \{f_{\theta, \eta} - f_{\theta, \eta_0} : \theta \in \Theta, \eta \in T_n\}$, which admits $2F_n$ as an envelope function. For every $\delta > 0$, it holds that $J(\delta, \mathcal{F}_n^0, \|\cdot\|_{2, P_{Q_0, g^r}}) \lesssim J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})$. Consequently, by (b) and (ii) in Lemma 2.12 below, there exists $\delta_0 > 0$ and $n_0 \geq 1$ such that, for all $n \geq n_0$, $J(\delta_0, \mathcal{F}_n^0, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq \alpha \varepsilon$. Define $T_n^0(\delta_0) \equiv \{\eta \in T_n : \sup_{\theta \in \Theta} P_{Q_0, g^r}(f_{\theta, \eta} - f_{\theta, \eta_0}) \leq \delta_0^2\}$, and $\mathcal{F}_n^0(\delta_0) \equiv \{f_{\theta, \eta} - f_{\theta, \eta_0} : \theta \in \Theta, \eta \in T_n^0(\delta_0)\} \subset \mathcal{F}_n^0$. By assumption, there exists $n_1 \geq 1$ such that $P(\eta_n \notin T_n^0(\delta_0)) \leq \varepsilon$ whenever $n \geq n_1$.

Set $n \geq \max(n_0, n_1)$. By the Markov inequality, and because $\eta_n \in T_n^0(\delta_0)$ implies $\widetilde{\mathcal{F}}_n^0 \subset \mathcal{F}_n^0(\delta_0)$, it holds that

$$\begin{aligned} & P\left(\sqrt{n} \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\widetilde{\mathcal{F}}_n^0} \geq \alpha\right) \\ & \leq P(\eta_n \notin T_n^0(\delta_0)) + \alpha^{-1} E\left(\sqrt{n} \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\widetilde{\mathcal{F}}_n^0} \mathbf{1}\{\eta_n \in T_n^0(\delta_0)\}\right) \end{aligned} \quad (2.50)$$

$$\begin{aligned} & \leq \varepsilon + \alpha^{-1} E\left(\sqrt{n} \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_n^0(\delta_0)} \mathbf{1}\{\eta_n \in T_n^0(\delta_0)\}\right) \\ & \leq \varepsilon + \alpha^{-1} E\left(\sqrt{n} \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_n^0(\delta_0)}\right). \end{aligned} \quad (2.51)$$

By Lemma 2.9, whose conditions are met,

$$\begin{aligned} & E\left(\sqrt{n} \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_n^0(\delta_0)}\right) \\ & \lesssim J(\delta_0, \mathcal{F}_n^0(\delta_0), \|\cdot\|_{2, P_{Q_0, g^r}}) + \sqrt{n} P_{Q_0, g^r} F_n \mathbf{1}\{F_n > \sqrt{n} a_n(\delta_0)/2\} \end{aligned} \quad (2.52)$$

$$\begin{aligned} & \lesssim J(\delta_0, \mathcal{F}_n^0, \|\cdot\|_{2, P_{Q_0, g^r}}) \\ & \quad + a_n(\delta_0)^{-1} P_{Q_0, g^r} F_n^2 \mathbf{1}\{F_n > \sqrt{n} a_n(\delta_0)/2\}, \end{aligned} \quad (2.53)$$

where $a_n(\delta_0) \equiv \delta_0 / \sqrt{\text{Log}N(\delta_0, \mathcal{F}_n^0(\delta_0), \|\cdot\|_{2, P_{Q_0, g^r}})}$. By (i) in Lemma 2.12 below, $m \mapsto J(\delta_0, \mathcal{F}_m^0(\delta_0), \|\cdot\|_{2, P_{Q_0, g^r}})$ is a bounded function. We also know that, for all $m \geq 1$,

$$J(\delta_0, \mathcal{F}_m^0(\delta_0), \|\cdot\|_{2, P_{Q_0, g^r}}) \geq \delta_0 \sqrt{\text{Log}N(\delta_0, \mathcal{F}_m^0(\delta_0), \|\cdot\|_{2, P_{Q_0, g^r}})} = \delta_0^2 / a_m(\delta_0).$$

In particular, $m \mapsto a_m(\delta_0)$ must be bounded away from 0. Let $c > 0$ be such that $a_m(\delta_0) \geq c$ for all $m \geq 1$. With this in mind, (2.53) implies

$$E \left(\sqrt{n} \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_n^0(\delta_0)} \right) \leq J(\delta_0, \mathcal{F}_n^0, \|\cdot\|_{2, P_{Q_0, g^r}}) + c^{-1} P_{Q_0, g^r} F_n^2 \mathbf{1}\{F_n > \sqrt{nc}/2\},$$

where $J(\delta_0, \mathcal{F}_n^0, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq \alpha \varepsilon$ by construction. Assumption (b) guarantees that there exists $n_2 \geq 1$ such that $m \geq n_2$ implies $P_{Q_0, g^r} F_m^2 \mathbf{1}\{F_m > \sqrt{mc}/2\} \leq \alpha c \varepsilon$. In summary, provided that $n \geq \max(n_0, n_1, n_2)$, (2.51) and (2.53) yield $P(\sqrt{n} \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_n^0} \geq \alpha) \leq 3\varepsilon$. In other words, $\sqrt{n} \|P_n - P_{Q_0, \mathbf{g}_n}\|_{\mathcal{F}_n^0} = o_P(1)$. This completes the proof. \square

The next two lemmas proved useful in our demonstrations.

Lemma 2.11. *Let \mathcal{F} be a uniformly bounded class of measurable, real-valued functions. Let h, h' be two measurable, bounded, real-valued functions. We do not assume that $h, h' \in \mathcal{F}$. Set $\delta > 0$.*

- Define \mathcal{F}' equal either to $\{f - h : f \in \mathcal{F}\}$, or $\{f|h| : f \in \mathcal{F}\}$, or $\{f|h| + f'|h'| : f, f' \in \mathcal{F}\}$, or $\{|f| : f \in \mathcal{F}\}$, or $\{f^2 : f \in \mathcal{F}\}$, or $\{\phi(f) : f \in \mathcal{F}\}$ where ϕ is non-decreasing and Lipschitz, or $\{h \log(f) + (1-h) \log(1-f) : f \in \mathcal{F}\}$ if the functions in \mathcal{F} and h take their values in $[0, 1]$ and are uniformly bounded away from 0 and 1. It holds that $J(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \lesssim J(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$.
- Define $\mathcal{F}' = \{\sqrt{f} : f \in \mathcal{F}\}$ if the functions in \mathcal{F} are non-negative. It holds that $J(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \lesssim J(\sqrt{\delta}, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$.

Proof of Lemma 2.11. Fix $\delta > 0$ and $M > 0$ such that $\sup_{f \in \mathcal{F}} \|f\|_\infty \leq M < \infty$. Let $N \equiv N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, and consider a collection of δ -brackets $\{[l_i, u_i] : i \leq N\}$ that covers \mathcal{F} .

- Case $\mathcal{F}' = \{f - h : f \in \mathcal{F}\}$. The collection of δ -brackets obtained by substituting $[l_i - h, u_i - h]$ for $[l_i, u_i]$, all $i \leq N$, covers \mathcal{F}' . This proves the first claim.
- Case $\mathcal{F}' = \{f|h| : f \in \mathcal{F}\}$. The collection of $\delta \|h\|_\infty$ -brackets obtained by substituting $[l_i|h|, u_i|h|]$ for $[l_i, u_i]$, all $i \leq N$, covers \mathcal{F}' . This proves the second claim.

- *Case $\mathcal{F}' = \{f|h + f'|h' : f, f' \in \mathcal{F}\}$.* The collection of brackets consisting of $[l_i|h + l_j|h', u_i|h + u_j|h']$, all $i, j \leq N$, covers \mathcal{F}' . Consider $i, j \leq N$, and set $\gamma_{ij} \equiv u_i|h + u_j|h'$, $\lambda_{ij} \equiv l_i|h + l_j|h'$, $c \equiv 2\sqrt{\|h\|_\infty^2 + \|h'\|_\infty^2}$: it holds that $P_{Q_0, g^r}(\gamma_{ij} - \lambda_{ij})^2 \leq c^2 \delta^2$.

Therefore, $N(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\delta/c, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})^2$, from which the third claim follows.

- *Case $\mathcal{F}' = \{|f| : f \in \mathcal{F}\}$.* Set $f \in \mathcal{F}$, and assume without loss of generality that $f \in [l_1, u_1]$. Define $F_+ \equiv \mathbf{1}\{f > 0\}$, $F_- \equiv \mathbf{1}\{f < 0\}$, $G_+ \equiv \mathbf{1}\{l_1 > 0\}$, $G_- \equiv \mathbf{1}\{u_1 < 0\}$, and $G_0 \equiv \mathbf{1}\{l_1 \leq 0 \leq u_1\}$. Then

$$F_+(l_1)_+ + F_-(u_1)_- \leq |f| \leq F_+u_1 - F_-l_1$$

with

$$F_+(l_1)_+ + F_-(u_1)_- = G_+l_1 - G_-u_1 \equiv \lambda_1,$$

and

$$\begin{aligned} F_+u_1 - F_-l_1 &= G_+u_1 - G_-l_1 + G_0(F_+u_1 - F_-l_1) \\ &\leq G_+u_1 - G_-l_1 + G_0(u_1 - l_1) \equiv \gamma_1. \end{aligned}$$

Thus $\lambda_1 \leq |f| \leq \gamma_1$, where $\gamma_1 - \lambda_1 = u_1 - l_1$, hence $P_{Q_0, g^r}(\gamma_1 - \lambda_1)^2 \leq \delta^2$.

Therefore, $N(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\delta, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, from which the fourth claim follows.

- *Case $\mathcal{F}' = \{f^2 : f \in \mathcal{F}\}$.* Set $f \in \mathcal{F}$, and assume without loss of generality that $f \in [l_1, u_1]$. Let $[\lambda_1, \gamma_1]$ be the bracket that we just built. The inequalities $\lambda_1 \geq 0$ and $f^2 \leq M^2$ imply that $\lambda_1^2 \leq f^2 \leq \min(\gamma_1^2, M^2)$. Set $\lambda_2 \equiv \lambda_1$, $\gamma_2 \equiv \sqrt{\min(\gamma_1^2, M^2)}$ so that $\lambda_2^2 \leq f^2 \leq \gamma_2^2$. Obviously, $\gamma_2^2 - \lambda_2^2 \leq 2\gamma_2(\gamma_2 - \lambda_2) \leq 2M(\gamma_1 - \lambda_1)$, hence $P_{Q_0, g^r}(\gamma_2^2 - \lambda_2^2)^2 \leq 4M^2 \delta^2$.

Therefore, $N(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\delta/2M, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, from which the fifth claim follows.

- *Case $\mathcal{F}' = \{\phi(f) : f \in \mathcal{F}\}$.* Say that ϕ is c -Lipschitz. The collection of $c\delta$ -brackets obtained by substituting $[\phi(l_i), \phi(u_i)]$ for $[l_i, u_i]$, all $i \leq N$, covers \mathcal{F}' . This proves the sixth claim.
- *Case $\mathcal{F}' = \{h \log(f) + (1-h) \log(1-f) : f \in \mathcal{F}\}$.* Set $f \in \mathcal{F}$, and assume without loss of generality that $f \in [l_1, u_1]$ and $0 < \inf_{f \in \mathcal{F}} f \leq l_1 \leq u_1 < \sup_{f \in \mathcal{F}} f < 1$. Define $\lambda_3 \equiv h \log(l) + (1-h) \log(1-l)$ and $\gamma_3 \equiv h \log(u) + (1-h) \log(1-u)$. It holds that

$\lambda_3 \leq h \log(f) + (1-h) \log(1-f) \leq \gamma_3$. Moreover, $0 \leq \gamma_3 - \lambda_3 \lesssim (u_1 - l_1)$ because \log is Lipschitz on any compact subset of $(0, 1)$. Consequently, there exists $c \geq 1$ such that $P_{Q_0, g^r}(\gamma_3 - \lambda_3)^2 \leq c^2 \delta^2$.

Therefore, $N(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\delta/c, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, from which the seventh claim follows.

- *Case $\mathcal{F}' = \{\sqrt{f} : f \in \mathcal{F}\}$.* Set $f \in \mathcal{F}$, and assume without loss of generality that $f \in [l_1, u_1]$ and $l_1 \geq 0$. Then $\sqrt{l_1} \leq \sqrt{f} \leq \sqrt{u_1}$. Moreover, $(\sqrt{u_1} - \sqrt{l_1})^2 \leq (\sqrt{u_1} - \sqrt{l_1})(\sqrt{u_1} + \sqrt{l_1}) = u_1 - l_1$. The Cauchy-Schwarz inequality yields $P_{Q_0, g^r}(u_1 - l_1) \leq \sqrt{P_{Q_0, g^r}(u_1 - l_1)^2} \leq \sqrt{\delta}$.

Therefore, $N(\delta, \mathcal{F}', \|\cdot\|_{2, P_{Q_0, g^r}}) \leq N(\sqrt{\delta}, \mathcal{F}, \|\cdot\|_{2, P_{Q_0, g^r}})$, from which the eighth claim follows.

This completes the proof. \square

Lemma 2.12. *For each $n \geq 1$, let \mathcal{F}_n be a class of measurable, real-valued functions such that $\delta_n = o(1)$ implies $J(\delta_n, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = o(1)$. Then (i) $J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = O(1)$ for every $\delta > 0$, and (ii) for every $\varepsilon > 0$, there exist $\delta > 0$ and $n_1 \geq 1$ such that $J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \leq \varepsilon$ for all $n \geq n_1$.*

Proof of Lemma 2.12. We prove (i) and (ii) by contradiction.

Suppose there exists $\delta > 0$ such that $\limsup_{n \rightarrow \infty} J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = \infty$. Without loss of generality, we can assume that $J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \geq 2^{2n}$ for each $n \geq 1$. Now,

$$J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = J(\delta/2, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) + \int_{\delta/2}^{\delta} \sqrt{\log N(\varepsilon, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})} d\varepsilon,$$

with

$$\begin{aligned} 2J(\delta/2, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) &\geq \delta \sqrt{\log N(\delta/2, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})} \\ &\geq 2 \int_{\delta/2}^{\delta} \sqrt{\log N(\varepsilon, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}})} d\varepsilon. \end{aligned}$$

Therefore, $J(\delta, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \geq 2^{2n}$ implies $J(\delta/2, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \geq 2^{2n}/2$ hence, by recursion,

$$J(\delta/2^n, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) \geq 2^{2n}/2^n = 2^n.$$

The sequence $\{\delta_n\}_{n \geq 1}$ given by $\delta_n = \delta/2^n$ satisfies $\delta_n = o(1)$ and $\lim_{n \rightarrow \infty} J(\delta_n, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) = \infty$, in contradiction with the assumption of the lemma. This completes the proof of (i).

Now, assume that there exists $\varepsilon > 0$ such that, for all $\delta > 0$, there exists $n_1 \geq 1$ for which $J(\delta, \mathcal{F}_{n_1}, \|\cdot\|_{2, P_{Q_0, g^r}}) > \varepsilon$. In particular, we can construct by recursion an increasing sequence $\{\varphi(n)\}_{n \geq 1}$ such that, for all $n \geq 1$, $J(1/n, \mathcal{F}_{\varphi(n)}, \|\cdot\|_{2, P_{Q_0, g^r}}) > \varepsilon$. This induces the existence of a sequence $\{\delta_n\}_{n \geq 1}$ such that $\delta_n = o(1)$ and $\limsup_{n \rightarrow \infty} J(\delta_n, \mathcal{F}_n, \|\cdot\|_{2, P_{Q_0, g^r}}) > \varepsilon$, in contradiction with the assumption of the lemma. This completes the proof of (ii). \square

Figure 2.1: **Performance TMLE across 500 simulations.** Each row corresponds to a performance measure (top: bias, middle: sample variance, bottom: MSE). Each column corresponds to a working model for the optimal randomization scheme (left: mis-specified working model \mathcal{G}_1^m , right: correctly specified working model \mathcal{G}_1^c). The red and green dots correspond to our CARA RCT with different working models for the conditional response (red: LASSO working model $\mathcal{Q}_{1,n}^\ell$, green: parametric working model $\mathcal{Q}_{1,n}^p$). The blue dots correspond to a RCT with a fixed design set to the balanced randomization scheme g^b and $\mathcal{Q}_{1,n}^p$ as (fixed) parametric working model for the conditional response.




Figure 2.2: **Empirical coverage of CLT-based 95% CIs across 500 simulations.** Each column corresponds to a working model for the optimal randomization scheme (left: mis-specified working model \mathcal{G}_1^m , right: correctly specified working model \mathcal{G}_1^c). The red and green dots correspond to our CARA RCT with different working models for the conditional response (red: LASSO working model $\mathcal{Q}_{1,n}^l$, green: parametric working model $\mathcal{Q}_{1,n}^p$). The blue dots correspond to a RCT with a fixed design set to the balanced randomization scheme g^b and $\mathcal{Q}_{1,n}^p$ as (fixed) parametric working model for the conditional response. The yellow lines indicate the confidence levels 95% and 94%.





Figure 2.3: **Hypotheses testing to assess, across 500 simulations, the quality of coverage guaranteed by the CLT-based 95%-CI.** Each row corresponds to a null hypothesis $H_0^{1-\alpha} : \pi \geq 1 - \alpha$ (top: $\alpha = 5\%$, bottom: $\alpha = 6\%$), where π is the actual coverage guaranteed by each CI, which should satisfy by construction $\pi \geq 95\%$. Each column corresponds to a working model for the optimal randomization scheme (left: mis-specified working model \mathcal{G}_1^m , right: correctly specified working model \mathcal{G}_1^c). The red and green colors correspond to our CARA RCT with different working models for the conditional response (red: LASSO working model $\mathcal{Q}_{1,n}^l$, green: parametric working model $\mathcal{Q}_{1,n}^p$). The blue color correspond to a RCT with a fixed design set to the balanced randomization scheme g^b and $\mathcal{Q}_{1,n}^p$ as (fixed) parametric working model for the conditional response. The yellow line indicates the threshold 0.05.



Chapter 3

Asymptotic Theory for Cross-validated Targeted Maximum Likelihood Estimation

3.1 Introduction

Current practice in statistics often involves fitting parametric or stringent semi-parametric regression models and using statistical inference for the regression coefficients in these models. These models are always wrong, and as a consequence the point estimates and confidence intervals are biased. Large sample sizes are not reducing this bias, but enhances false rejections of null hypotheses. In addition, this parametric approach does not focus on carefully translating the scientific question of interest in terms of a target parameter of the probability distribution of the data.

In van der Laan and Rubin (2006) we introduced targeted maximum likelihood estimation (TMLE) in semiparametric models, which incorporates adaptive estimation (e.g., loss based super learning) of the relevant part of the data generating distribution, and subsequently carries out a targeted bias reduction by maximizing the log-likelihood (or other loss function for the relevant part) of a "clever" parametric working-model through the initial estimator, treating the initial estimator as off-set, and possibly iterates this targeted updating step till convergence. The target parameter of the resulting updated estimator is then evaluated, and is called the TMLE of the target parameter of the data generating distribution. This estimator is, by definition, a substitution estimator, and, under regularity conditions, is a double robust semiparametric efficient estimator. We refer the reader to van der Laan, Rose, and Gruber (September, 2009) for applications of TMLE.

The use of adaptive estimators raises the question till what degree we can still rely

on the central limit theorem for statistical inference. Our previous theorems show that under empirical process conditions and rate of convergence conditions, one can indeed still prove asymptotic linearity, and thereby obtain CLT-based inference. The empirical process conditions puts some bounds on how adaptive the initial estimator can be. Indeed, we have experienced that using as initial estimator an adaptive regression algorithm that overfits the data such as the machine learning algorithm Random Forest can negatively impact the bias reduction performance of the subsequent TMLE-step. In this paper we present a version of targeted MLE that uses V-fold sample splitting. We refer to this as the cross-validated targeted MLE (CV-TMLE). We formally establish its asymptotics under stated conditions that avoid such empirical process conditions. The implications of this theorem for the role of super learning (i.e., adaptive estimation) in construction of semiparametric efficient estimators of target parameters is discussed. We also present a direct application of this version of targeted MLE to the estimation of the additive causal effect of a binary treatment on an outcome. We shall see that under mild conditions (e.g. initial estimators need not be consistent), the resulting estimator is of the form

$$\psi_n^* - \psi_0 = (P_n - P_0)IC(P_0) + R_n,$$

where the remainder is second order. The conditions for asymptotic linearity of ψ_n^* thus follow from the analysis of this second order term.

The organization of this article is as follows. In section 2 we formally present the TMLE using V-fold sample splitting for the initial estimator (CV-TMLE). In section 3 we focus on the one-step CV-TMLE and present a theorem establishing its asymptotics. The conditions and implications of the theorem are discussed. We also present an extension of the theorem with more practical implications. In section 4 the theorem is demonstrated for the cross-validated TMLE of the causal effect of a binary treatment on a continuous or binary outcome. We discuss the implications of the theorem in strategies for estimating the target parameter of the data generating distribution using data adaptive estimators combined with CV-TMLE. In section 5 we present a theorem for the general iterative CV-TMLE, and its conditions are discussed. We end this article with a discussion. Technical derivations are put in the Appendix.

3.2 The TMLE using V-fold sample splitting for initial estimator.

Let $O \sim P_0$ and the probability distribution P_0 is known to be an element of a statistical model \mathcal{M} . We observe n i.i.d. copies O_1, \dots, O_n of O and wish to estimate a particular multivariate target parameter $\Psi(P_0)$. Let P_n denote the empirical probability distribution

of O_1, \dots, O_n so that estimators can be represented as mappings from an empirical distribution to the parameter space of the parameter it is estimating: for example, $P_n \rightarrow \hat{\Psi}(P_n)$ denotes an estimator of $\psi_0 = \Psi(P_0)$.

We assume that Ψ is pathwise differentiable along a class of 1-dimensional sub-models $\{P_h(\varepsilon) : \varepsilon\}$ indexed by a choice h in an index set \mathcal{H} : i.e., there exists a fixed d -variate function $D(P) = (D_1(P), \dots, D_d(P))$ so that for all $h \in \mathcal{H}$

$$\left. \frac{d}{d\varepsilon} \Psi(P_h(\varepsilon)) \right|_{\varepsilon=0} = PD(P)S(h),$$

where $S(h)$ is the score of $\{P_h(\varepsilon) : \varepsilon\}$ at $\varepsilon = 0$. Here we used the notation $PS = \int S(o)dP(o)$ for the expectation of a function S of O .

We assume that a parameter $Q : \mathcal{M} \rightarrow \mathcal{Q}$ is chosen so that $\Psi(P_0) = \Psi^1(Q(P_0))$. For convenience, we will refer to both mappings with Ψ , so we will abuse notation by using interchangeably $\Psi(Q(P))$ as well as $\Psi(P)$. Let $g : \mathcal{M} \rightarrow \mathcal{G}$ be so that for all $P \in \mathcal{M}$,

$$D^*(P) = D^*(Q(P), g(P)).$$

In other words, the canonical gradient only depends on P through a relevant part $Q(P)$ of P and a nuisance parameter $g(P)$ of P .

Let $\mathcal{L}^\infty(K)$ be the class of functions of O with bounded supremum norm over a set of K so that $P_0(O \in K) = 1$, endowed with the supremum norm. We assume there exists a uniformly bounded loss function $L : \mathcal{Q} \rightarrow \mathcal{L}^\infty(K)$ so that

$$Q(P_0) = \arg \min_{Q \in \mathcal{Q}} P_0 L(Q),$$

where, we remind the reader that $P_0 L(Q) = \int L(Q)(o)dP_0(o)$. In addition, we assume that for each $P \in \mathcal{M}$, for a specified d -dimensional (hardest) parametric model $\{P(\varepsilon) : \varepsilon\} \subset \mathcal{M}$ through P at $\varepsilon = 0$ and with score $D^*(P)$ at $\varepsilon = 0$,

$$\left\langle \left. \frac{d}{d\varepsilon} L(Q(P(\varepsilon))) \right|_{\varepsilon=0} \right\rangle \supset \langle D^*(P) \rangle.$$

We are now ready to define a targeted maximum likelihood estimator. Let $P_n \rightarrow \hat{Q}(P_n)$ be an initial estimator of $Q_0 = Q(P_0)$. Let $P_n \rightarrow \hat{g}(P_n)$ be an initial estimator of $g_0 = g(P_0)$. Given \hat{Q}, \hat{g} , let $P_n \rightarrow \hat{Q}(P_n)(\varepsilon)$ be a family of estimators indexed by ε chosen so that

$$\left\langle \left. \frac{d}{d\varepsilon} L(\hat{Q}(P_n)(\varepsilon)) \right|_{\varepsilon=0} \right\rangle \supset \langle D^*(\hat{Q}(P_n), \hat{g}(P_n)) \rangle. \quad (3.1)$$

Here we used the notation $\langle h \rangle$ for the linear span spanned by the components of $h = (h_1, \dots, h_k)$. One can think of $\{\hat{Q}(P_n)(\varepsilon) : \varepsilon\} \subset \mathcal{M}$ as a submodel through $\hat{Q}(P_n)$ with

parameter ε , chosen so that the derivative(or score) at $\varepsilon = 0$ yields a function that equals or spans the efficient influence curve at the initial estimator $(\hat{Q}(P_n), \hat{g}(P_n))$. Note that this submodel for fluctuating $\hat{Q}(P_n)$ uses the estimator $\hat{g}(P_n)$ in its definition.

Let $B_n \in \{0, 1\}^n$ be a random vector indicating a split of $\{1, \dots, n\}$ into a training and validation sample: $\mathcal{T} = \{i : B_n(i) = 0\}$ and $\mathcal{V} = \{i : B_n(i) = 1\}$. Let P_{n,B_n}^0, P_{n,B_n}^1 be the empirical probability distributions of the training and validation sample, respectively. For a given cross-validation scheme $B_n \in \{0, 1\}^n$, we now define

$$\varepsilon_n^0 = \hat{\varepsilon}(P_n) \equiv \arg \min_{\varepsilon} E_{B_n} P_{n,B_n}^1 L(\hat{Q}(P_{n,B_n}^0)(\varepsilon)).$$

This now yields an update $\hat{Q}(P_{n,B_n}^0)(\varepsilon_n^0)$ of $\hat{Q}(P_{n,B_n}^0)$ for each split B_n .

As a side-note, it is of interest to point out that this cross-validated selector of ε equals the cross-validation selector among the library of candidate estimators $P_n \rightarrow \hat{Q}(P_n)(\varepsilon)$ of Q_0 indexed by ε . As a consequence, we can apply the results for the cross-validation selector that show that it is asymptotically equivalent with the so called oracle selector. Formally, consider the oracle selector

$$\tilde{\varepsilon}_n^0 \equiv \arg \min_{\varepsilon} E_{B_n} P_0 L(\hat{Q}(P_{n,B_n}^0)(\varepsilon)).$$

If, in addition to uniform boundedness, we assume that the loss function also satisfies

$$M_2 = \sup_{Q \in \mathcal{Q}} \frac{\text{VAR}\{L(Q) - L(Q_0)\}}{E_0\{L(Q) - L(Q_0)\}} < \infty,$$

then the results in van der Laan and Dudoit (2003) and van der Vaart, Dudoit, and van der Laan (2006) imply that we have the following finite sample inequality:

$$\begin{aligned} 0 &\leq EE_{B_n} P_0 \{L(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n^0)) - L(\hat{Q}(P_{n,B_n}^0)(\tilde{\varepsilon}_n^0))\} \\ &\leq 2\sqrt{c} \frac{1}{\sqrt{n}} \sqrt{EE_{B_n} P_0 \{L(\hat{Q}(P_{n,B_n}^0)(\tilde{\varepsilon}_n^0)) - L(Q_0)\}}. \end{aligned}$$

Here c can be explicitly bounded by M_2 and an upper bound of L . This finite sample inequality gives us insight in the benefit of using cross-validation to select the amount of fluctuation ε , since it shows that ε_n^0 will be close to the oracle selector $\tilde{\varepsilon}_n^0$ for any choice of initial estimators (even if the initial estimator is extremely data adaptive).

One could now iterate this updating process of the training sample specific estimators: define $\hat{Q}^1(P_{n,B_n}^0) = \hat{Q}(P_{n,B_n}^0)(\varepsilon_n^0)$, define the family of fluctuations $P_n \rightarrow \hat{Q}^1(P_n)(\varepsilon)$ satisfying the derivative condition (3.1), and set

$$\varepsilon_n^1 = \arg \min_{\varepsilon} E_{B_n} P_{n,B_n}^1 L(\hat{Q}^1(P_{n,B_n}^0)(\varepsilon)),$$

resulting in another update $\hat{Q}^1(P_{n,B_n}^0)(\varepsilon_n^1)$ for each B_n . This process is iterated till $\varepsilon_n^k = 0$ (or close enough to zero). The final update will be denoted with $\hat{Q}^*(P_{n,B_n}^0)$ for each split B_n . The targeted MLE is now defined as

$$\hat{\Psi}(P_n) \equiv E_{B_n} \Psi(\hat{Q}^*(P_{n,B_n}^0)).$$

We refer to this as the *cross-validated TMLE* (CV-TMLE).

In a variety of examples, the convergence occurs in one step (i.e., $\varepsilon_n^1 = 0$ already). In this case, we write $\varepsilon_n \equiv \varepsilon_n^0$ and

$$\hat{\Psi}(P_n) = E_{B_n} \Psi(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n)).$$

Cross-validated TMLE when one of the components is linear in data generating distribution

The CV-TMLE presented above can be generalized to the case where only one component of the initial estimator $\hat{Q}(P_n)$ should be updated using a parametric working fluctuation model, while the other component can be estimated using a substitution estimator plugging in the empirical probability distribution function (i.e., an NPMLE). In this case, it is not necessary to target the second component since it is already an unbiased estimator. Formally, consider a decomposition of Q into (Q_1, Q_2) , such that $Q_2 \rightarrow \Psi(Q_1, Q_2)$ is linear, and $Q_2(P)$ is linear in P itself so that it is sensible to estimate it with an empirical probability distribution. Suppose that the canonical gradient D^* can be decomposed as

$$D^*(P) = D_1^*(P) + D_2^*(P),$$

where $D_1^*(P_0)$ is the canonical gradient of the map

$$P \rightarrow \Psi(Q_1(P), Q_2(P_0))$$

at $P = P_0$. Assume also that $D_1^*(P)$ does not depend on $Q_2(P)$.

Under these assumptions we can apply the CV-TMLE algorithm to obtain a targeted estimator of $Q_1(P_0)$, while not updating the initial estimator of $Q_2(P_0)$. In this case, the parametric fluctuation model satisfies

$$\left\langle \frac{d}{d\varepsilon} L(\hat{Q}_1(P_n)(\varepsilon)) \Big|_{\varepsilon=0} \right\rangle \supset \langle D_1^*(\hat{Q}_1(P_n), \hat{g}(P_n)) \rangle,$$

where $L()$ is now a loss function for $Q_1(P_0)$ only. For a given cross-validation scheme $B_n \in \{0, 1\}^n$, we define

$$\varepsilon_n^0 = \hat{\varepsilon}(P_n) \equiv \arg \min_{\varepsilon} E_{B_n} P_{n,B_n}^1 L(\hat{Q}_1(P_{n,B_n}^0)(\varepsilon)).$$

This now yields an update $\hat{Q}_1(P_{n,B_n}^0)(\varepsilon_n^0)$ of $\hat{Q}_1(P_{n,B_n}^0)$ for each split B_n . One could now iterate this updating process of the training sample specific estimators: define $\hat{Q}_1^1(P_{n,B_n}^0) = \hat{Q}_1(P_{n,B_n}^0)(\varepsilon_n^0)$, define the family of fluctuations $P_n \rightarrow \hat{Q}_1^1(P_n)(\varepsilon)$ satisfying the derivative condition (3.1), and set

$$\varepsilon_n^1 = \arg \min_{\varepsilon} E_{B_n} P_{n,B_n}^1 L(\hat{Q}_1^1(P_{n,B_n}^0)(\varepsilon)),$$

resulting in another update $\hat{Q}_1^1(P_{n,B_n}^0)(\varepsilon_n^1)$ for each B_n . This process is iterated till $\varepsilon_n^k = 0$ (or close enough to zero). The final update will be denoted with $\hat{Q}_1^*(P_{n,B_n}^0)$ for each split B_n . The resulting CV-TMLE of ψ_0 is given by

$$\hat{\Psi}(P_n) = E_{B_n} \Psi(\hat{Q}_1^*(P_{n,B_n}^0), \hat{Q}_2(P_{n,B_n}^1)).$$

We will illustrate this estimator with an application to the additive causal effect of a binary treatment on a continuous or binary outcome in section 4.

3.3 Asymptotics for the one-step cross-validated TMLE

In this section we analyze the cross-validated targeted MLE that converge in one step. The theorem carries relevance in general since it establishes the theoretical behavior of the targeted MLE updating algorithm. For convenience, in this section and the next ε_n^0 is simply denoted with ε_n . In the following theorem, convergence in probability always refers to convergence when n converges to infinity.

Definition 3.1. *For a class of functions, \mathcal{F} , whose elements are functions f that map \mathcal{O} into a real number, we define the entropy integral*

$$\text{Entro}(\mathcal{F}) \equiv \int_0^\infty \sqrt{\log \sup_Q N(\varepsilon \| F \|_{Q,2}, \mathcal{F}, L_2(Q))} d\varepsilon,$$

where $N(\varepsilon, \mathcal{F}, L_2(Q))$ is the covering number, defined as the minimal number of balls of radius $\varepsilon > 0$ needed to cover \mathcal{F} , using the $L_2(Q)$ -norm when defining a ball of radius ε . In addition, F is defined as the envelope of \mathcal{F} which is a function F so that $|f| \leq F$ for all $f \in \mathcal{F}$.

We refer to van der Vaart and Wellner (1996b) for empirical process theory. We state the following lemma (Lemma 2.14.1 in van der Vaart and Wellner (1996b)) for ease of reference.

Lemma 3.1. Let \mathcal{F} denote a class of measurable functions of O . Let $G_n = \sqrt{n}(P_n - P_0)$. Then

$$E \left(\sup_{f \in \mathcal{F}} |G_n f| \right) \leq \text{Entro}(\mathcal{F}) \sqrt{P_0 F^2}.$$

The following result is an immediate application of lemma 3.1.

Lemma 3.2. Suppose $\|\varepsilon_n - \varepsilon_0\| \xrightarrow{P} 0$. For each sample split of B_n , we condition on P_{n,B_n}^0 and consider a class of measurable functions of O :

$$\mathcal{F}(P_{n,B_n}^0) \equiv \{f_\varepsilon(P_{n,B_n}^0) \equiv f(\varepsilon, P_{n,B_n}^0) - f(\varepsilon_0, P_0) : \varepsilon\},$$

where the index set contains ε_n with probability tending to 1. For a deterministic sequence $\delta_n \rightarrow 0$, define the subclasses

$$\mathcal{F}_{\delta_n}(P_{n,B_n}^0) \equiv \{f_\varepsilon \in \mathcal{F}(P_{n,B_n}^0) : \|\varepsilon - \varepsilon_0\| < \delta_n\}.$$

If for deterministic sequence $\delta_n \rightarrow 0$, we have

$$E \left\{ \text{Entro}(\mathcal{F}_{\delta_n}(P_{n,B_n}^0)) \sqrt{P_0 F(\delta_n, P_{n,B_n}^0)^2} \right\} \rightarrow 0 \text{ as } n \rightarrow \infty,$$

where $F(\delta_n, P_{n,B_n}^0)$ is the envelope of $\mathcal{F}_{\delta_n}(P_{n,B_n}^0)$, then

$$\sqrt{n}(P_{n,B_n}^1 - P_0) \{f(\varepsilon_n, P_{n,B_n}^0) - f(\varepsilon_0, P_0)\} = o_P(1).$$

Theorem 3.1. Let $\hat{Q}(P_n)$, $\hat{g}(P_n)$ be an initial estimator of Q_0 , g_0 , respectively. In the following, $\hat{Q}(P_0)$ and $\hat{g}(P_0)$ denote the limits of these estimators, not necessarily equal to Q_0 and g_0 , respectively.

Uniformly bounded loss function: We assume that $\{\hat{Q}(P_n)(\varepsilon) : \varepsilon\} \in \mathcal{Q}$ with probability 1, the loss function $L(Q)$ for Q_0 is uniformly bounded in $Q \in \mathcal{Q}$, and over a support of $O \sim P_0$:

$$M_1 = \sup_Q \sup_O |L(Q)(O)| < \infty.$$

Let $B_n \in \{0, 1\}^n$ be a random vector indicating a split of $\{1, \dots, n\}$ into a training and validation sample. Suppose B_n is uniformly distributed over a finite support.

Consider the estimator defined above

$$\hat{\Psi}(P_n) = E_{B_n} \Psi(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n)).$$

If the parameter $P \rightarrow \Psi(Q(P))$ satisfies

A1:

$$\Psi(Q(P)) - \Psi(Q_0) = -P_0 D^*(Q(P), g_0) + O_P(\|\Psi(Q(P)) - \Psi(Q_0)\|^2).$$

Then

$$\begin{aligned}
\hat{\Psi}(P_n) - \psi_0 &= E_{B_n} (P_{n,B_n}^1 - P_0) D^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \\
&+ E_{B_n} P_0 \{ D^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), g_0) \} \\
&- E_{B_n} P_0 \{ D^* (Q_0, \hat{g}(P_{n,B_n}^0)) - D^* (Q_0, g_0) \} \\
&+ O_P(\|\hat{\Psi}(P_n) - \psi_0\|^2). \tag{3.2}
\end{aligned}$$

Consider $\varepsilon_0 = \varepsilon(P_0)$ such that $\|\varepsilon_n - \varepsilon_0\| \xrightarrow{P} 0$. Suppose the following assumption also holds:

A2: (Given $\|\varepsilon_n - \varepsilon_0\| \xrightarrow{P} 0$)

For each sample split B_n , condition on P_{n,B_n}^0 and define the class of functions

$$\mathcal{F}(P_{n,B_n}^0) \equiv \{O \rightarrow D^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon), \hat{g}(P_{n,B_n}^0)) - D^* (\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) : \varepsilon\},$$

where the set over which ε varies is chosen so that it contains ε_n with probability tending to 1. In addition, for a deterministic sequence δ_n converging to zero as $n \rightarrow \infty$, we also define the sequence of sub-classes

$$\mathcal{F}_{\delta_n}(P_{n,B_n}^0) \equiv \{f_\varepsilon \in \mathcal{F}(P_{n,B_n}^0) : \|\varepsilon - \varepsilon_0\| < \delta_n\}.$$

Assume that for deterministic sequence δ_n converging to 0, we have

$$E \text{Entro}(\mathcal{F}_{\delta_n}(P_{n,B_n}^0)) \sqrt{P_0 F^2(\delta_n, P_{n,B_n}^0)} \rightarrow 0 \text{ as } n \rightarrow \infty,$$

where $F(\delta_n, P_{n,B_n}^0)$ is the envelope of $\mathcal{F}_{\delta_n}(P_{n,B_n}^0)$.

Then we have:

$$\begin{aligned}
\hat{\Psi}(P_n) - \psi_0 &= (P_n - P_0) D^* (\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) + o_P(1/\sqrt{n}) \\
&+ E_{B_n} P_0 \{ D^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), g_0) \} \\
&- E_{B_n} P_0 \{ D^* (Q_0, \hat{g}(P_{n,B_n}^0)) - D^* (Q_0, g_0) \} \\
&+ O_P(\|\hat{\Psi}(P_n) - \psi_0\|^2). \tag{3.3}
\end{aligned}$$

Furthermore, suppose $\hat{g}(P_n) = g_0$, that is g_0 is known, as in the case of an RCT. Then

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) D^* (\hat{Q}(P_0)(\varepsilon_0), g_0) + o_P(1/\sqrt{n}). \tag{3.4}$$

If, in addition to $\hat{g}(P_n) = g_0$, we also have $\hat{Q}(P_0)(\varepsilon_0) = Q_0$, then $\hat{\Psi}(P_n)$ is in fact asymptotically efficient:

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0)D^*(Q_0, g_0) + o_P(1/\sqrt{n}). \quad (3.5)$$

More generally, suppose $\hat{g}(P_0) = g_0$. Let \tilde{Q} denote the limit of $\hat{Q}(P_n)(\varepsilon_n)$ which is not necessarily Q_0 . Assume in addition

A3:

$$\begin{aligned} & E_{B_n} P_0 \{ D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), g_0) \} \\ & - E_{B_n} P_0 \{ D^*(\tilde{Q}, \hat{g}(P_{n,B_n}^0)) - D^*(\tilde{Q}, g_0) \} \\ & = o_P(1/\sqrt{n}). \end{aligned}$$

A4: For some mean zero function $IC'(P_0) \in L_0^2(P_0)$, we have

$$\begin{aligned} & E_{B_n} P_0 \{ D^*(\tilde{Q}, \hat{g}(P_{n,B_n}^0)) - D^*(\tilde{Q}, g_0) \} \\ & - E_{B_n} P_0 \{ D^*(Q_0, \hat{g}(P_{n,B_n}^0)) - D^*(Q_0, g_0) \} \\ & = (P_n - P_0)IC'(P_0) + o_P(1/\sqrt{n}). \end{aligned}$$

NOTE: If $\hat{Q}(P_n)(\varepsilon_n)$ converges to Q_0 then A4 is automatically true with $IC' \equiv 0$.

Then $\hat{\Psi}(P_n)$ is asymptotically linear

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) \{ D^*(\hat{Q}(P_0)(\varepsilon_0), g_0) + IC'(P_0) \} + o_P(1/\sqrt{n}).$$

Note that the choice of the initial estimator of the Q -function affects the update estimator $\hat{Q}(P_{n,B_n}^0)(\varepsilon_n)$ and its subsequent limit. In our expansion in (3.3), this choice would play out in the influence curve $D^*(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0))$, which is a function of this limit, and hence would affect the efficiency bound, as well as the quadratic residual term $O_P(\|\hat{\Psi}(P_n) - \psi_0\|^2)$. The TMLE update will not alter consistency of the Q -function itself, only the consistency of the target parameter estimate, which is a function of this Q -function.

Proof of Theorem 3.1:

From definition of ε_n and the one-step convergence of $\hat{Q}(P)(\varepsilon_n)$, we have that

$$E_{B_n} P_{n,B_n}^1 D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) = 0.$$

The double robustness of D^* guarantees $P_0 D^*(Q_0, g) = 0$ for all g . Combining this result with A1, we readily have (3.2):

$$\hat{\Psi}(P_n) - \psi_0 = E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \quad (3.6)$$

$$+ E_{B_n} P_0 \{ D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), g_0) \} \quad (3.7)$$

$$- E_{B_n} P_0 \{ D^*(Q_0, \hat{g}(P_{n,B_n}^0)) - D^*(Q_0, g_0) \} \quad (3.8)$$

$$+ O_P(\| \hat{\Psi}(P_n) - \psi_0 \|^2).$$

We may rewrite (3.6) as

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \\ &= E_{B_n} (P_{n,B_n}^1 - P_0) \{ D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) \} \\ &+ E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)). \end{aligned}$$

An application of lemma 3.2 and A2 implies that for each sample split B_n ,

$$\begin{aligned} & (P_{n,B_n}^1 - P_0) \{ D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) \} \\ &= o_P(1/\sqrt{n}). \end{aligned}$$

Since B_n is uniformly distributed on a finite support, it now follows that indeed

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) \{ D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) \} \\ &= o_P(1/\sqrt{n}). \end{aligned}$$

In other words, the term (3.6) is given by

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \\ &= E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) + o_P(1/\sqrt{n}). \end{aligned}$$

This result and the established equality in (3.2) now prove (3.3).

Now, if $\hat{g}(P_n) = g_0$, then the (3.7) and (3.8) are exactly 0. Consequently, (3.3) becomes

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= (P_n - P_0) D^*(\hat{Q}(P_0)(\varepsilon_0), g_0) \\ &+ o_P(1/\sqrt{n}) + O_P(\| \hat{\Psi}(P_n) - \psi_0 \|^2). \end{aligned}$$

However, taking $\| \cdot \|$ on both sides of the equality above yields $\| \hat{\Psi}(P_n) - \psi_0 \| = o_P(1/\sqrt{n})$. We thereby have asymptotically linearity of $\hat{\Psi}(P_n)$:

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) D^*(\hat{Q}(P_0)(\varepsilon_0), g_0) + o_P(1/\sqrt{n}).$$

If, in addition, $\hat{Q}(P_0)(\varepsilon_0) = Q_0$, then the influence curve is indeed the efficient influence curve $D^*(Q_0, g_0)$.

Next we consider a more general case where $\hat{g}(P_0) = g_0$. Let \tilde{Q} be the limit of $\hat{Q}(P_n)(\varepsilon_n)$. It is not necessarily the case that $\tilde{Q} = Q_0$. We now rewrite the established equality (3.3) to account for \tilde{Q} :

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= (P_n - P_0) D^* (\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) + o_P(1/\sqrt{n}) \\ &+ E_{B_n} P_0 \{ D^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), g_0) \} \\ &- E_{B_n} P_0 \{ D^* (\tilde{Q}, \hat{g}(P_{n,B_n}^0)) - D^* (\tilde{Q}, g_0) \} \\ &+ E_{B_n} P_0 \{ D^* (\tilde{Q}, \hat{g}(P_{n,B_n}^0)) - D^* (\tilde{Q}, g_0) \} \\ &- E_{B_n} P_0 \{ D^* (Q_0, \hat{g}(P_{n,B_n}^0)) - D^* (Q_0, g_0) \} \\ &+ O_P(\| \hat{\Psi}(P_n) - \psi_0 \|^2). \end{aligned}$$

From A3, the term

$$\begin{aligned} &E_{B_n} P_0 \{ D^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), g_0) \} \\ &- E_{B_n} P_0 \{ D^* (\tilde{Q}, \hat{g}(P_{n,B_n}^0)) - D^* (\tilde{Q}, g_0) \} \\ &= o_P(1/\sqrt{n}). \end{aligned}$$

From A4, the term

$$\begin{aligned} &E_{B_n} P_0 \{ D^* (\tilde{Q}, \hat{g}(P_{n,B_n}^0)) - D^* (\tilde{Q}, g_0) \} \\ &- E_{B_n} P_0 \{ D^* (Q_0, \hat{g}(P_{n,B_n}^0)) - D^* (Q_0, g_0) \} \\ &= (P_n - P_0) IC'(P_0) + o_P(1/\sqrt{n}). \end{aligned}$$

Therefore (3.3) becomes

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= (P_n - P_0) \{ D^* (\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) + IC'(P_0) \} + o_P(1/\sqrt{n}) \\ &+ O_P(\| \hat{\Psi}(P_n) - \psi_0 \|^2). \end{aligned}$$

Taking $\| \cdot \|$ on both sides again yields $\| \hat{\Psi}(P_n) - \psi_0 \| = o_P(1/\sqrt{n})$. We thereby have the desired result

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) \{ D^* (\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) + IC'(P_0) \} + o_P(1/\sqrt{n}).$$

□

Remarks about conditions of Theorem 3.1

To understand assumption A1 we note the following. By general property of the efficient influence curve, we have

$$\Psi(P) - \Psi(P_0) = -P_0 D^*(P) + R(P, P_0),$$

where the specifics of the behavior of the remainder R as a function of P, P_0 depend on the particular data structure, semiparametric model, and target parameter. For example, for linear parameters on convex models we have $\Psi(P) - \Psi(P_0) = -P_0 D^*(P)$ exact, as shown in van der Laan (2006).

Under no conditions on the estimators, we determined an exact identity (3.2) for the cross-validated TMLE minus its target ψ_0 , which already provides the main insights about the performance of this estimator. It shows that the analysis of the CV-TMLE involves a cross-validated empirical process term applied to the efficient influence curve, and a remainder term (In many examples we shall see that this remainder is second order). The cross-validated empirical process term is nice because it involves, for each sample split, an empirical mean over a validation sample of an estimated efficient influence curve that is largely estimated based on the training sample. Based on this, one would predict that one can establish a CLT for this cross-validated empirical process term without having to enforce restrictive entropy conditions on the support of (i.e., class of functions that contains) the estimated efficient influence curve (and thereby limit the adaptiveness of the initial estimators). This is formalized by A2 and our second result (3.3), which replaces the cross-validated empirical process term by an empirical mean of mean zero random variables $D^*(\hat{Q}(P_0)(\epsilon_0), \hat{g}(P_0))$ plus a negligible $o_P(1/\sqrt{n})$ -term. This result only requires the positivity assumption, and *that the estimators converge to a target*. That is, under essentially no conditions beyond the positivity assumption, the CV-TMLE minus the true ψ_0 , behaves as an empirical mean of mean zero i.i.d. random variables (which thus converges to a normal distribution, by CLT), plus a specified remainder term. In particular, we control bias of the estimator by making this remainder term as small as possible.

Regarding assumption A2 we note the following. Combined with lemma 3.2, A2 implies that the cross-validated empirical process term minus an empirical mean of mean zero random variables converges to 0 at root-n rate. The entropy-term in A2 concerns the entropy of a class of functions that are indexed by a finite dimensional parameter. Such entropies are bounded under very weak conditions, mainly that the class of functions are uniformly bounded. As a consequence, to obtain the wished convergence, one first simply provides a bound on the entropy of $\mathcal{F}(P)$ for a fixed P uniformly in all P . In this way, it remains to show that

$$EP_0 \mathbf{F}^2(\delta_n, P_{n, B_n}^0) \rightarrow 0 \text{ as } n \rightarrow \infty.$$

In other words, one shows that the $L^2(P_0)$ -norm of the envelope converges to zero for $\delta_n \rightarrow 0$ and P_{n,B_n}^0 converging to P_0 . Again, this is mainly a consistency condition on $\hat{Q}(P_n)(\varepsilon_n)$ (with respect to its limit, which is not necessarily Q_0). More importantly, we do not require that the entropy of the space of initial estimator $\hat{Q}(P_n)$, and thereby also the entropy of $\hat{g}(P_n)$, is controlled. The latter are typical conditions putting strong restrictions on how data adaptive the estimators \hat{Q} and \hat{g} can be, but these conditions are now completely avoided. This result allows us to fully utilize data adaptive estimators to make the remainder term negligible.

Moreover, in an RCT g_0 is known, and one might set $\hat{g}(P_n) = g_0$, so that the remainder term is exactly equal to zero, giving us the asymptotic linearity (3.4) of the CV-TMLE under no other conditions than the positivity assumption and convergence of $\hat{Q}(P_n)$ to some fixed function. This teaches us the remarkable lesson that in an RCT, one can use very aggressive super learning without causing any violations of the conditions, and one will approximate the efficiency bound at smaller sample sizes than otherwise. In particular, in an RCT in which we use a consistent estimator \hat{Q} the CV-TMLE is asymptotically efficient, as stated in (3.5). That is, in an RCT, this theorem teaches us that CV-TMLE with adaptive estimation of \bar{Q}_0 is the way to go.

In more general types of studies, when $\hat{g}(P_n) \neq g_0$, the remainder may not be exactly zero. But its form, as described in (3.3), will allow us to identify the necessary conditions and general strategies for estimation of Q_0 and g_0 to make this term negligible. We will illustrate this in our example with estimation of additive causal effect of binary treatment on an outcome.

Implication for the use of super learning The importance of using super learning for estimation of both Q_0 and g_0 is now clear. Super learning is essential to make the remainder as small as possible, for controlling bias. Interestingly, at least asymptotically, there seems to be no price for using super learning, but only benefits: one wants the remainder term in (3.3) to be small, and that requires approximating the true Q_0 and g_0 well, and simultaneously, the use of very data adaptive estimators did not affect the conditions required for the analysis of the asymptotically linear term in (3.3), due to the V-fold sample splitting. Therefore, to control the bias term asymptotically, the utilization of super learning is essential, while it also improves the efficiency of the first order term. Further investigation of the required conditions for the bias-term will have to teach us if there will be any trade-off between obtaining a good rate of convergence and the entropy of the estimators. We will return to this issue in our example and its following remarks.

Asymptotics for CV-TMLE when one of the components is linear in data generating distribution

We now study the asymptotics of the CV-TMLE described in section 3.2 when the algorithm converges in one step.

Consider a decomposition of Q into $Q = (Q_1, Q_2)$, such that $Q_2 \mapsto \Psi(Q_1, Q_2)$ is linear, and $Q_2(P)$ is linear in P itself. Suppose we can decompose the canonical gradient D^* as

$$\begin{aligned} D^*(Q_1(P), Q_2(P), g(P)) &= D_1^*(Q_1(P), g(P)) \\ &+ D_2^*(Q_1(P), g(P)) + D_3^*(Q_1(P), Q_2(P), g(P)), \end{aligned}$$

where $D_1^*(P_0)$ is the canonical gradient of the map

$$P \mapsto \Psi(Q_1(P), Q_2(P_0))$$

at $P = P_0$. In our additive causal effect example in next section, Q_2 plays the role of the marginal distribution of the baseline covariates and $Q_1 = E(Y|1, W) - E(Q|0, W)$. Since $\Psi(Q_0)$ only involves taking an average w.r.t. the covariate distribution, $Q_{2,0}$ is naturally estimated with its empirical distribution. In our example, $D_2^*(Q_1, g) = Q_1$ and $D_3^*(Q_1, Q_2, g) = -\Psi(Q_1, Q_2)$.

Under certain conditions on D_2^* and D_3^* , the asymptotic results of previous theorem extend naturally to the CV-TMLE where $Q_{1,0}$ is estimated using a fluctuation model and $Q_{2,0}$ is estimated using a substitution estimator plugging in the empirical distribution.

Theorem 3.2. *Consider a decomposition of Q into $Q = (Q_1, Q_2)$, such that $Q_2 \mapsto \Psi(Q_1, Q_2)$ is linear and $Q_2(P)$ is linear in P .*

Suppose the canonical gradient D^ can be decomposed into*

$$\begin{aligned} D^*(Q_1(P), Q_2(P), g(P)) &= D_1^*(Q_1(P), g(P)) \\ &+ D_2^*(Q_1(P), g(P)) + D_3^*(Q_1(P), Q_2(P), g(P)), \end{aligned}$$

where $D_1^*(P_0)$ is the canonical gradient of the map

$$P \mapsto \Psi(Q_1(P), Q_2(P_0))$$

at $P = P_0$. Denote $D'^* \equiv (D_1^* + D_2^*)$

Let $\hat{Q}_1(P_n)$, $\hat{Q}_2(P_n)$, $\hat{g}(P_n)$ be estimators of $Q_{1,0}$, $Q_{2,0}$, g_0 , respectively. We will denote their limits with $\hat{Q}_1(P_0)$, $\hat{Q}_2(P_0)$, and $\hat{g}(P_0)$, which are not necessarily equal to $Q_{1,0}$, $Q_{2,0}$ and g_0 , respectively.

Uniformly bounded loss function: We assume that $\{\hat{Q}_1(P_n)(\varepsilon) : \varepsilon\} \in \mathcal{Q}$ with probability

1, the loss function $L(Q_1)$ for $Q_{1,0}$ is uniformly bounded in $Q_1 \in \mathcal{Q}$, and over a support of $O \sim P_0$:

$$M_1 = \sup_Q \sup_O |L(Q_1)(O)| < \infty.$$

Let $B_n \in \{0, 1\}^n$ be a random vector indicating a split of $\{1, \dots, n\}$ into a training and validation sample. Suppose B_n is uniformly distributed over a finite support.

Denote $\hat{Q}(P_n, B_n)(\varepsilon_n) \equiv (\hat{Q}_1(P_{n, B_n}^0)(\varepsilon_n), \hat{Q}_2(P_{n, B_n}^1))$, and let

$$\hat{\Psi}(P_n) \equiv E_{B_n} \Psi(\hat{Q}_1(P_{n, B_n}^0)(\varepsilon_n), \hat{Q}_2(P_{n, B_n}^1)).$$

If the parameter $P \rightarrow \Psi(Q(P))$ satisfies

A1:

$$\Psi(Q(P)) - \Psi(Q_0) = -P_0 D^*(Q(P), g_0) + O_P(\|\Psi(Q(P)) - \Psi(Q_0)\|^2),$$

and

A2:

$$\begin{aligned} & E_{B_n} P_{n, B_n}^1 D_2^*(\hat{Q}_1(P_{n, B_n}^0)(\varepsilon_n), \hat{g}(P_{n, B_n}^0)) \\ & + E_{B_n} P_{n, B_n}^1 D_3^*(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n, B_n}^0)) \\ & = 0. \end{aligned}$$

Then

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= E_{B_n} (P_{n, B_n}^1 - P_0) D^*(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n, B_n}^0)) \\ & + E_{B_n} P_0 \{D^*(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n, B_n}^0)) - D^*(\hat{Q}(P_n, B_n)(\varepsilon_n), g_0)\} \\ & - E_{B_n} P_0 \{D^*(Q_0, \hat{g}(P_{n, B_n}^0)) - D^*(Q_0, g_0)\} \\ & + O_P(\|\hat{\Psi}(P_n) - \psi_0\|^2), \end{aligned} \tag{3.9}$$

Let $\varepsilon_0 = \varepsilon(P_0)$ be such that $\|\varepsilon_n - \varepsilon_0\| \xrightarrow{P} 0$.

Suppose the following assumptions also hold

A3: For each sample split B_n

$$\begin{aligned} & \sqrt{n}(P_{n, B_n}^1 - P_0) \{D_3^*(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n, B_n}^0)) - D_3^*(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0))\} \\ & = o_P(1), \end{aligned}$$

where $\hat{Q}(P_0)(\varepsilon_0) = (\hat{Q}_1(P_0)(\varepsilon_0), \hat{Q}_2(P_0))$.

A4: (Given $\|\varepsilon_n - \varepsilon_0\| \xrightarrow{P} 0$.)

Conditional on each P_{n,B_n}^0 , define the class of functions

$$\mathcal{F}(P_{n,B_n}^0) \equiv \{O \rightarrow D^*(\hat{Q}_1(P_{n,B_n}^0)(\varepsilon), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}_1(P_0)(\varepsilon_0), \hat{g}(P_0)) : \varepsilon\},$$

where the set over which ε varies is chosen so that it contains with probability tending to 1 ε_n . In addition, for a deterministic sequence δ_n converging to zero as $n \rightarrow \infty$, we also define the sequence of sub-classes

$$\mathcal{F}_{\delta_n}(P_{n,B_n}^0) \equiv \{f_\varepsilon \in \mathcal{F}(P_{n,B_n}^0) : \|\varepsilon - \varepsilon_0\| < \delta_n\}.$$

Assume that for deterministic sequence δ_n converging to 0, we have

$$E \text{Entr}(\mathcal{F}_{\delta_n}(P_{n,B_n}^0)) \sqrt{P_0 F(\delta_n, P_{n,B_n}^0)^2} \rightarrow 0 \text{ as } n \rightarrow \infty,$$

where $F(\delta_n, P_{n,B_n}^0)$ is the envelope of $\mathcal{F}_{\delta_n}(P_{n,B_n}^0)$.

Then we have:

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= (P_n - P_0) D^*(\hat{Q}_1(P_0)(\varepsilon_0), \hat{Q}_2(P_0), \hat{g}(P_0)) + o_P(1/\sqrt{n}) \\ &+ E_{B_n} P_0 \{D^*(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_n, B_n)(\varepsilon_n), g_0)\} \\ &- E_{B_n} P_0 \{D^*(Q_0, \hat{g}(P_{n,B_n}^0)) - D^*(Q_0, g_0)\} \\ &+ O_P(\|\hat{\Psi}(P_n) - \psi_0\|^2), \end{aligned} \quad (3.10)$$

Furthermore, suppose $\hat{g}(P_n) = g_0$. Then

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) D^*(\hat{Q}_1(P_0)(\varepsilon_0), \hat{Q}_2(P_0), g_0) + o_P(1/\sqrt{n}).$$

If, in addition to $\hat{g}(P_n) = g_0$, we also have $\hat{Q}_1(P_0)(\varepsilon_0) = Q_{1,0}$ and $\hat{Q}_2(P_0) = Q_{2,0}$, then $\hat{\Psi}(P_n)$ is in fact asymptotically efficient

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) D^*(Q_0, g_0) + o_P(1/\sqrt{n}).$$

More generally, suppose $\hat{g}(P_0) = g_0$. Let \tilde{Q}_1 denote the limit of $\hat{Q}_1(P_n)(\varepsilon_n)$ which is not necessarily $Q_{1,0}$. Assume in addition

A5:

$$\begin{aligned} &E_{B_n} P_0 \{D^*(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_n, B_n)(\varepsilon_n), g_0)\} \\ &- E_{B_n} P_0 \{D^*(\tilde{Q}_1, \hat{Q}_2(P_0), \hat{g}(P_{n,B_n}^0)) - D^*(\tilde{Q}_1, \hat{Q}_2(P_0), g_0)\} \\ &= o_P(1/\sqrt{n}). \end{aligned}$$

A6: For some mean zero function $IC'(P_0) \in L_0^2(P_0)$, we have

$$\begin{aligned} & E_{B_n} P_0 \{ D^*(\tilde{Q}_1, \hat{Q}_2(P_0), \hat{g}(P_{n,B_n}^0)) - D^*(\tilde{Q}_1, \hat{Q}_2(P_0), g_0) \} \\ & - E_{B_n} P_0 \{ D^*(Q_0, \hat{g}(P_{n,B_n}^0)) - D^*(Q_0, g_0) \} \\ & = (P_n - P_0) IC'(P_0) + o_P(1/\sqrt{n}). \end{aligned}$$

NOTE: If $\hat{Q}_1(P_n)(\varepsilon_n)$ converges to $Q_{1,0}$ and $\hat{Q}_2(P_n)$ converges to $Q_{2,0}$ then A6 is automatically true with $IC' \equiv 0$.

Then $\hat{\Psi}(P_n)$ is asymptotically linear

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) \{ D^*(\hat{Q}_1(P_0)(\varepsilon_0), \hat{Q}_2(P_0), g_0) + IC'(P_0) \} + o_P(1/\sqrt{n}).$$

Proof of Theorem 3.2:

From definition of ε_n and one-step convergence of $\hat{Q}_1(P)(\varepsilon_n)$, we have that

$$E_{B_n} P_{n,B_n}^1 D_1^*(\hat{Q}_1(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) = 0.$$

Combining this result with A1, A2 and the double robustness of D^* , which guarantees $P_0 D^*(Q_0, g) = 0$ for all g , we readily have (3.9):

$$\hat{\Psi}(P_n) - \psi_0 = E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_{n,B_n})(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \quad (3.11)$$

$$+ E_{B_n} P_0 \{ D^*(\hat{Q}(P_{n,B_n})(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_{n,B_n})(\varepsilon_n), g_0) \} \quad (3.12)$$

$$- E_{B_n} P_0 \{ D^*(Q_0, \hat{g}(P_{n,B_n}^0)) - D^*(Q_0, g_0) \} \quad (3.13)$$

$$+ O_P(\| \hat{\Psi}(P_n) - \psi_0 \|^2),$$

On the other hand, we may rewrite (3.11) as

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \\ & = E_{B_n} (P_{n,B_n}^1 - P_0) \left\{ D^{*'}(\hat{Q}_1(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^{*'}(\hat{Q}_1(P_0)(\varepsilon_0), \hat{g}(P_0)) \right\} \\ & + (P_n - P_0) D^{*'}(\hat{Q}_1(P_0)(\varepsilon_0), \hat{g}(P_0)) \\ & + E_{B_n} (P_{n,B_n}^1 - P_0) \left\{ D_3^*(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D_3^*(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) \right\} \\ & + (P_n - P_0) D_3^*(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)) \end{aligned}$$

Applying the lemma 3.2 with A4 we have that

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) \left\{ D^{*'}(\hat{Q}_1(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^{*'}(\hat{Q}_1(P_0)(\varepsilon_0), \hat{g}(P_0)) \right\} \\ & = o_P(1/\sqrt{n}). \end{aligned}$$

It follows from this result and A3 that the term (3.11) becomes

$$\begin{aligned}
& E_{B_n} (P_{n,B_n}^1 - P_0) D^* (\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \\
&= E_{B_n} (P_{n,B_n}^1 - P_0) D^{*'} (\hat{Q}_1(P_0)(\varepsilon_0), \hat{g}(P_0)) \\
&+ (P_n - P_0) D_3^* (\hat{Q}_1(P_0)(\varepsilon_0), \hat{Q}_2(P_0), \hat{g}(P_0)) + o_P(1/\sqrt{n}) \\
&= E_{B_n} (P_{n,B_n}^1 - P_0) D^* (\hat{Q}_1(P_0)(\varepsilon_0), \hat{Q}_2(P_0), \hat{g}(P_0)) + o_P(1/\sqrt{n})
\end{aligned}$$

These results and the established equality in (3.9) now prove (3.10).

Similar steps as in the proof of Theorem 3.1 now complete this proof. \square

Remark on conditions of Theorem 3.2

For some parameters, it is more efficacious to only target one component of Q_0 while estimating the other component using a substitution estimator plugging in the empirical distribution. Theorem 3.2 teaches us that the resulting CV-TMLE, under this partial-targeting scheme, has all the desired properties of its full-targeting counterpart. The analysis of the theoretical behavior of $\hat{\Psi}$ in Theorem 3.1 can be extended natural to obtain the results in Theorem 3.2 if D_2^* and D_3^* satisfy A2 and A3. These two conditions give us insight into when it is sensible to use this partial-targeting CV-TMLE for Q .

Condition A2 implies that one may still solve the estimation equation by only targeting Q_1 and estimating Q_2 using the validation set, i.e.

$$E_{B_n} P_{n,B_n}^1 D^* (\hat{Q}_1(P_{n,B_n}^0)(\varepsilon_n), \hat{Q}_2(P_{n,B_n}^1), \hat{g}(P_{n,B_n}^0)) = 0.$$

This suggests that it's sensible to employ this partial-targeting scheme only if the estimator $(\hat{Q}_1(P_{n,B_n}^0)(\varepsilon_n), \hat{Q}_2(P_{n,B_n}^1))$ will be as good as its full-targeting counterpart in terms of solving the cross validated estimating equation.

In our examples, $D_3^*(Q_1, Q_2, g) = -\Psi(Q_1, Q_2)$, in which case A3 is automatically true since $(P_n - P_0)D_3^*(Q_1, Q_2, g) = 0$ for all Q_1, Q_2, g . In such instances, no requirements are imposed on the estimators, and thus the partial-targeting scheme is highly effective. However, when that is not the case, A3 implies that one will need to control the entropy of the class of estimators \hat{Q}_2 , since they will be evaluated at the training set P_{n,B_n}^1 . In these cases, the partial-targeting scheme may not be as effective as the full-targeting one.

3.4 Application of Theorem 3.2 to estimation of additive causal effect in nonparametric model

Let $O = (W, A, Y)$, W be a vector of baseline covariates, A a binary treatment variable, and Y an outcome of interest. Let \mathcal{M} be the class of all probability distributions for O . We

consider the parameter $\Psi : \mathcal{M} \rightarrow \mathbf{R}$

$$\Psi(Q(P)) = E_P [E_P(Y|W, A = 1) - E_P(Y|W, A = 0)].$$

Several estimators, in addition to TMLE, have been proposed for the estimation of this parameter: the G-comp estimator (Robins (1986)), the IPTW estimator (Hernan, Brumback, and Robins (2000); Robins (1999)), the DR-IPTW estimator (Robins and Rotnitzky (2001); Robins (2000); Robins, Rotnitzky, and van der Laan (2000)). We refer to van der Laan et al. (September, 2009), Gruber and van der Laan (2010), Stitelman and van der Laan (2010), and Petersen, Porter, S.Gruber, Wang, and van der Laan (2010) for comparisons of performance between TMLE and these various estimators.

Let $Q(P) = (\bar{Q}(P), Q_W(P))$, where $\bar{Q}(P)(W, A) \equiv E_P(Y|W, A)$ and $Q_W(P)$ is the density of the marginal probability distribution of W . For convenience, we will use $\bar{Q}(P)(W)$ to denote $E_P(Y|W, A = 1) - E_P(Y|W, A = 0)$. The distinctions will be clear from the arguments given to the function or from context. Let $g(P)(A|W) \equiv Pr_P(A|W)$. We also adopt the notations $\bar{Q}_0 \equiv \bar{Q}(P_0)$, $Q_{W,0} \equiv Q_W(P_0)$, and $g_0 \equiv g(P_0)$.

Our parameter of interest is Ψ evaluated at the distribution $P_0 \in \mathcal{M}$ of the observed O :

$$\psi_0 \equiv \Psi(Q_0) = E_{W,0} [E_0(Y|W, A = 1) - E_0(Y|W, A = 0)].$$

The canonical gradient of Ψ at $P \in \mathcal{M}$ is

$$\begin{aligned} D^*(Q(P), g(P))(O) &= \left\{ H_{g(P)}^*(A, W) (Y - \bar{Q}(P)(A, W)) \right\} \\ &+ \left\{ \bar{Q}(P)(W) - Q_W(P)\bar{Q}(P) \right\} \\ &\equiv D_Y^*(\bar{Q}(P), g(P)) + D_W^*(\bar{Q}(P), Q_W(P)), \end{aligned}$$

where

$$H_g^*(A, W) = \left(\frac{A}{g(1|W)} - \frac{1-A}{g(0|W)} \right).$$

For convenience, we will also use the notation

$$H_g^*(W) \equiv H_g^*(1, W) - H_g^*(0, W).$$

Firstly, note that the map $Q_W \mapsto \Psi(\bar{Q}, Q_W)$ is linear and $Q_W(P)$ is linear in P . Secondly, $D_Y^*(\bar{Q}_0, g_0)$ is the canonical gradient of the map $P \mapsto \Psi(\bar{Q}(P), Q_W(P_0))$ at $P = P_0$, and does not depend on $Q_W(P_0)$. In the following we present a TMLE of Q_0 where only the initial estimator $\hat{Q}(P_n)$ of \bar{Q}_0 is updated using a parametric working model $\hat{Q}(P_n)(\varepsilon)$, while the marginal distribution of W is estimated with the empirical distribution which is not updated. Given an appropriate loss function $L(\bar{Q})$ and initial estimators \hat{Q} and \hat{g} of \bar{Q}_0

and g_0 , respectively, the parametric working model $\{\hat{Q}(P_n)(\varepsilon) : \varepsilon\}$ will be selected such that

$$\frac{d}{d\varepsilon} L(\hat{Q}(P_n)(\varepsilon)) \Big|_{\varepsilon=0} = D_Y^*(\hat{Q}(P_n), \hat{g}(P_n)).$$

We consider here two possible loss functions for binary outcome or continuous outcomes $Y \in [0, 1]$.

Squared error loss function: The squared error loss function is given by

$$L(\bar{Q})(O) \equiv (Y - \bar{Q}(A, W))^2,$$

with the parametric working model

$$\hat{Q}(P_n)(\varepsilon) = \hat{Q}(P_n) + \varepsilon H_{\hat{g}(P_n)}^*.$$

Quasi-log-likelihood loss function: The quasi-log-likelihood loss function is given by

$$L(\bar{Q})(O) \equiv - (Y \log(\bar{Q}(W, A)) + (1 - Y) \log(1 - \bar{Q}(W, A))),$$

with parametric working model

$$\hat{Q}(P_n)(\varepsilon) = \frac{1}{1 + e^{-\text{logit}(\hat{Q}(P_n)) - \varepsilon H_{\hat{g}(P_n)}^*}}.$$

We note that we would use this loss function if Y is binary or Y is continuous with values in $(0, 1)$. If Y is a bounded continuous random variable with values in (a, b) , then we can still use this loss function by using the transformed outcome $Y^* = (Y - a)/(b - a)$ and mapping the obtained TMLE of the additive treatment effect on Y^* (and confidence intervals) into a TMLE of the additive treatment effect on Y (and confidence intervals).

It is important to point out that the TMLE of \bar{Q}_0 corresponding with both fluctuation models will converge in one step, since the clever covariate $H_{\hat{g}(P_n)}^*$ in the update of \hat{Q} does not involve \hat{Q} .

Let $B_n \in \{0, 1\}^n$ be a random vector indicating a split of $\{1, \dots, n\}$ into a training and validation sample: $\mathcal{T} = \{i : B_n(i) = 0\}$ and $\mathcal{V} = \{i : B_n(i) = 1\}$. Let $P_{n, B_n}^0, P_{n, B_n}^1$ be the empirical probability distributions of the training and validation sample, respectively. Given the parametric working model, the optimal ε_n is selected using cross validation:

$$\varepsilon_n = \arg \min_{\varepsilon} E_{B_n} P_{n, B_n}^1 L(\hat{Q}(P_{n, B_n}^0)(\varepsilon)).$$

In particular, the one-step convergence implies that ε_n satisfies

$$0 = E_{B_n} P_{n,B_n}^1 D_Y^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)). \quad (3.14)$$

At each sample split B_n , we define the TMLE of Q_0 at (P_n, B_n) as

$$\hat{Q}(P_n, B_n)(\varepsilon_n) \equiv \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{Q}_W(P_{n,B_n}^1) \right).$$

The TMLE of ψ_0 is defined as

$$\hat{\Psi}(P_n) \equiv E_{B_n} \Psi \left(\hat{Q}(P_n, B_n)(\varepsilon_n) \right) = E_{B_n} \Psi \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{Q}_W(P_{n,B_n}^1) \right).$$

Next, we illustrate the theoretical advantages of this estimator under both loss functions. We will show that under a natural rate condition on the initial estimators \hat{Q} and \hat{g} , the resulting TMLE $\hat{\Psi}(P_n)$ is asymptotically linear, and when \hat{g} and \hat{Q} are consistent, its influence curve is indeed the efficient influence curve.

Squared error loss for \bar{Q}

Let the loss function for \bar{Q}_0 be:

$$L(\bar{Q})(O) \equiv (Y - \bar{Q}(A, W))^2,$$

and consider the parametric working model through $\bar{Q}(P)$ for any $P \in \mathcal{M}$:

$$\bar{Q}(P)(\varepsilon) = \bar{Q}(P) + \varepsilon H_{\bar{g}(P)}^*.$$

Then, for given initial estimators \hat{g} and \hat{Q} , we have

$$\hat{Q}(P_n)(\varepsilon) = \hat{Q}(P_n) + \varepsilon H_{\hat{g}(P_n)}^*. \quad (3.15)$$

The cross validation selector of ε in (3.15) is defined as

$$\begin{aligned} \varepsilon_n &\equiv \arg \min_{\varepsilon} E_{B_n} P_{n,B_n}^1 L \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon) \right) \\ &= \arg \min_{\varepsilon} E_{B_n} \sum_{i, B_n(i)=1} \left(Y_i - \hat{Q}(P_{n,B_n}^0)(\varepsilon)(A_i, W_i) \right)^2. \end{aligned}$$

At each sample split B_n , we define the TMLE of Q_0 at (P_n, B_n) as

$$\hat{Q}(P_n, B_n)(\varepsilon_n) \equiv \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{Q}_W(P_{n,B_n}^1) \right),$$

where $\hat{Q}_W(P_{n,B_n}^1)$ is the marginal empirical distribution of W in the validation set. The TMLE of ψ_0 is defined as

$$\hat{\Psi}(P_n) \equiv E_{B_n} \Psi \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{Q}_W(P_{n,B_n}^1) \right).$$

We will now apply the main Theorem 3.2 to $\hat{\Psi}(P_n)$ which provides us with the following result.

Theorem 3.3. *Consider the setting above under the squared error loss function.*

Let $B_n \in \{0, 1\}^n$ be a random vector indicating a split of $\{1, \dots, n\}$ into a training and validation sample. Suppose B_n is uniformly distributed on a finite support.

Let \hat{Q} and \hat{g} be initial estimators of \bar{Q}_0 and g_0 . In the following, $\hat{Q}(P_0)$ and $\hat{g}(P_0)$ denote limits of these estimators, not necessarily equal to \bar{Q}_0 and g_0 , respectively.

The cross-validated TMLE satisfies

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= E_{B_n} (P_{n,B_n}^1 - P_0) D^* \left(\hat{Q}(P_{n,B_n})(\varepsilon_n), \hat{g}(P_{n,B_n}^0) \right) \\ &+ E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) (g_0 - \hat{g}(P_{n,B_n}^0)) \right\}. \end{aligned} \quad (3.2)$$

Suppose now that there exists a constant $L > 0$ such that $P_0(|Y| < L) = 1$.

Consider the following definition:

$$\varepsilon_0 \equiv \arg \min_{\varepsilon} P_0 L(\hat{Q}(P_0)(\varepsilon)).$$

Suppose that this minimum exists and satisfies the derivative equation

$$0 = P_0 D_Y(P_0, \varepsilon_0),$$

where

$$\begin{aligned} D_Y(P, \varepsilon) &\equiv \frac{d}{d\varepsilon} L(\hat{Q}(P)(\varepsilon))(O) \\ &= \left(Y - \hat{Q}(P)(A, W) - \varepsilon H_{\hat{g}(P)}^*(A, W) \right) H_{\hat{g}(P)}^*(A, W) \\ &= D_Y^* \left(\hat{Q}(P)(\varepsilon), \hat{g}(P) \right). \end{aligned}$$

If there are multiple minima, then it is assumed that the argmin is uniquely defined and selects one of these minima.

Suppose that \hat{Q} and \hat{g} satisfy the following conditions:

1. There exists a closed bounded set $K \subset \mathbb{R}^k$ containing ε_0 such that ε_n belongs to K with probability 1;
2. For some $\delta > 0$, $P(1 - \delta > \hat{g}(P_n)(1|W) > \delta) = 1$;
3. For some $K > 0$, $P(|\hat{Q}(P_n)(A, W)| < K) = 1$;
4.
$$\int_W (\hat{g}(P_n)(1|W) - \hat{g}(P_0)(1|W))^2 dQ_{W,0}(w) \rightarrow 0 \text{ in probability};$$
5. For $a = 0, 1$,

$$\int_W \left(\hat{Q}(P_n)(a, w) - \hat{Q}(P_0)(a, w) \right)^2 dQ_{W,0}(w) \rightarrow 0 \text{ in probability.}$$

Then,

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= (P_n - P_0) \left\{ D_Y^* \left(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0) \right) + \hat{Q}(P_0)(\varepsilon_0) \right\} \\ &+ E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) (g_0 - \hat{g}(P_{n,B_n}^0)) \right\} \\ &+ o_P(1/\sqrt{n}). \end{aligned} \quad (3.3)$$

Furthermore, If $\hat{g}(P_n) = g_0$, the TMLE estimator $\hat{\Psi}(P_n)$ is asymptotically linear estimator of ψ_0 :

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) D^* \left(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0) \right) + o_P(1/\sqrt{n}), \quad (3.4)$$

where $\hat{Q}(P_0)(\varepsilon_0) = (\hat{Q}(P_0)(\varepsilon_0), Q_{W,0})$.

If, in addition to $\hat{g}(P_n) = g_0$, $\hat{Q}(P_0) = \bar{Q}_0$, which implies that $\hat{Q}(P_0)(\varepsilon_0) = \bar{Q}_0$, then $\hat{\Psi}(P_n)$ is an asymptotically efficient estimator of ψ_0 :

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) D^*(Q_0, g_0) + o_P(1/\sqrt{n}). \quad (3.5)$$

More generally, if the limits satisfy $\hat{g}(P_0) = g_0$ and $\hat{Q}(P_0) = \bar{Q}_0$, and if the convergence satisfies

$$\begin{aligned} &\sqrt{E_{B_n} P_0 \left(\frac{g_0 - \hat{g}(P_{n,B_n}^0)}{g_0 \hat{g}(P_{n,B_n}^0)} \right)^2} \sqrt{E_{B_n} P_0 \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \bar{Q}_0 \right)^2} \\ &= o_P(1/\sqrt{n}), \end{aligned} \quad (3.16)$$

then $\hat{\Psi}(P_n)$ is an asymptotically efficient estimator of ψ_0 :

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) D^*(Q_0, g_0) + o_P(1/\sqrt{n}).$$

Consider now the case that $\hat{g}(P_0) = g_0$, but $\hat{Q}(P_0) \neq \bar{Q}_0$. If the convergence satisfies

$$\begin{aligned} & \sqrt{E_{B_n} P_0 \left(\frac{g_0 - \hat{g}(P_{n,B_n}^0)}{g_0 \hat{g}(P_{n,B_n}^0)} \right)^2} \sqrt{E_{B_n} P_0 \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(P_0)(\varepsilon_0) \right)^2} \\ &= o_P(1/\sqrt{n}), \end{aligned} \quad (3.17)$$

and $P_0 \left\{ H_{\hat{g}(P_n)}^* \left(\hat{Q}(P_0)(\varepsilon_0) - \bar{Q}_0 \right) \right\}$ is an asymptotically linear estimator of the function $P_0 \left\{ H_{\hat{g}(P_0)}^* \left(\hat{Q}(P_0)(\varepsilon_0) - \bar{Q}_0 \right) \right\}$ with influence curve IC' , then $\hat{\Psi}(P_n)$ is an asymptotically linear estimator of ψ_0 :

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) \left\{ D^*(\hat{Q}(P_0)(\varepsilon_0), g_0) + IC' \right\} + o_P(1/\sqrt{n}).$$

For convenience of reference, we state several simple but useful results in the proof of the theorem.

Lemma 3.3. *If X_n converges to X in probability, and there exists $\eta > 0$ such that $P(|X_n| < \eta) = 1$, then $E|X_n - X|^r \rightarrow 0$ for $r \geq 1$.*

Lemma 3.4. *Suppose \hat{g} is such that for some $\delta > 0$, $P(1 - \delta > \hat{g}(P_n)(1 | W) > \delta) = 1$. If for $a = 0, 1$, \hat{g} satisfies $P_{W,0}(\hat{g}(P_n) - \hat{g}(P_0))^2 \xrightarrow{P} 0$, then we have that $P_0 \left(H_{\hat{g}(P_n)}^* - H_{\hat{g}(P_0)}^* \right)^4$, $P_0 \left(H_{\hat{g}(P_n)}^* - H_{\hat{g}(P_0)}^* \right)^2$, $P_0 \left(H_{\hat{g}(P_n)}^* - H_{\hat{g}(P_0)}^* \right)$ and $P_0 \left((H_{\hat{g}(P_n)}^*)^2 - (H_{\hat{g}(P_0)}^*)^2 \right)$ also converge to zero in probability.*

Lemma 3.5. *Suppose \hat{g} and \hat{Q} satisfy the conditions 2-5 in Theorem 3.3. Then, for each split B_n , for any $r \geq 1$,*

1. $EP_0 \left(\hat{Q}(P_{n,B_n}^0) H_{\hat{g}(P_{n,B_n}^0)}^* - \hat{Q}(P_0) H_{\hat{g}(P_0)}^* \right)^r \rightarrow 0$;
2. $EP_0 \left((Y - \hat{Q}(P_{n,B_n}^0)) H_{\hat{g}(P_{n,B_n}^0)}^* - (Y - \hat{Q}(P_0)) H_{\hat{g}(P_0)}^* \right)^r \rightarrow 0$;
3. $EP_0 \left((H_{\hat{g}(P_{n,B_n}^0)}^*)^2 - (H_{\hat{g}(P_0)}^*)^2 \right)^r \rightarrow 0$;
4. $EP_0 \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right)^r \rightarrow 0$.

We are now ready to prove Theorem 3.3.

Proof. Firstly, we wish to establish that

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= E_{B_n} (P_{n,B_n}^1 - P_0) D^* \left(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0) \right) \\ &+ E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) (g_0 - \hat{g}(P_{n,B_n}^0)) \right\}, \end{aligned}$$

where $\hat{Q}(P_n, B_n)(\varepsilon_n) = \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{Q}_W(P_{n,B_n}^1) \right)$.

Note that

$$\begin{aligned} &-P_0 D^*(Q(P), g_0) \\ &\equiv -P_0 \left\{ (Y - \bar{Q}(P)) H_{g_0}^* + \bar{Q}(P) - P_W(P) \bar{Q}(P) \right\} \\ &= - \left\{ P_0 Y H_{g_0}^* - P_0 \bar{Q}(P) H_{g_0}^* + P_{W,0} \bar{Q}(P) - P_W(P) \bar{Q}(P) \right\} \\ &= P_W(P) \bar{Q}(P) - P_0 Y H_{g_0}^* \\ &= \Psi(Q(P)) - \Psi(Q_0). \end{aligned}$$

Applying this result to each sample split of B_n and averaging, it follows that

$$\hat{\Psi}(P_n) - \psi_0 \equiv E_{B_n} \Psi \left(\hat{Q}(P_n, B_n)(\varepsilon_n) \right) - \Psi(Q(P_0)) = -E_{B_n} P_0 D^* \left(\hat{Q}(P_n, B_n)(\varepsilon_n), g_0 \right). \quad (3.18)$$

On the other hand,

$$\begin{aligned} &E_{B_n} P_{n,B_n}^1 D_W^* \left(\hat{Q}_W(P_{n,B_n}^1), \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) \\ &\equiv E_{B_n} P_{n,B_n}^1 \left\{ \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - P_W(P_{n,B_n}^1) \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right\} \\ &= E_{B_n} \left\{ P_W(P_{n,B_n}^1) \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - P_W(P_{n,B_n}^1) \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right\} = 0. \end{aligned}$$

Moreover, it follows from the definition of ε_n and the one-step convergence of the chosen fluctuation model that $\left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0) \right)$ satisfies (3.14). Therefore, we have

$$\begin{aligned} &E_{B_n} P_{n,B_n}^1 D^* \left(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0) \right) \\ &\equiv E_{B_n} P_{n,B_n}^1 D_Y^* \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0) \right) \\ &+ E_{B_n} P_{n,B_n}^1 D_W^* \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{Q}_W(P_{n,B_n}^1) \right) \\ &= 0. \end{aligned} \quad (3.19)$$

Combining (3.18), (3.19) and robustness of D^* , $P_0 D^*(Q_0, g) = 0$ for all g , we may now rewrite $\hat{\Psi}(P_n) - \psi_0$ as

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= E_{B_n} (P_{n,B_n}^1 - P_0) D^* (\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \\ &+ E_{B_n} P_0 \{ D^* (\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^* (\hat{Q}(P_n, B_n)(\varepsilon_n), g_0) \} \\ &- E_{B_n} P_0 \{ D^* (Q_0, \hat{g}(P_{n,B_n}^0)) - D^* (Q_0, g_0) \}. \end{aligned}$$

The last two summands in this equality can be combined as

$$\begin{aligned} &E_{B_n} P_0 \{ D^* (\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) - D^* (\hat{Q}(P_n, B_n)(\varepsilon_n), g_0) \} \\ &- E_{B_n} P_0 \{ D^* (Q_0, \hat{g}(P_{n,B_n}^0)) - D^* (Q_0, g_0) \} \\ &\equiv E_{B_n} P_0 \left\{ D_Y^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) + D_W^* (\hat{Q}(P_n, B_n)(\varepsilon_n)) \right\} \\ &- E_{B_n} P_0 \left\{ D_Y^* (\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), g_0) + D_W^* (\hat{Q}(P_n, B_n)(\varepsilon_n)) \right\} \\ &- E_{B_n} P_0 \left\{ D_Y^* (\bar{Q}_0, \hat{g}(P_{n,B_n}^0)) + D_W^* (Q_0) \right\} \\ &+ E_{B_n} P_0 \left\{ D_Y^* (\bar{Q}_0, g_0) + D_W^* (Q_0) \right\} \\ &= E_{B_n} P_0 \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{g_0}^* \right) \\ &= E_{B_n} P_0 \left\{ \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) (-1)^{1+A} \frac{(g_0 - \hat{g}(P_{n,B_n}^0))}{g_0 \hat{g}(P_{n,B_n}^0)} \right\}. \end{aligned}$$

Therefore, we indeed have the desired expression (3.2):

$$\hat{\Psi}(P_n) - \psi_0 = E_{B_n} (P_{n,B_n}^1 - P_0) D^* (\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \quad (3.20)$$

$$+ E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) (g_0 - \hat{g}(P_{n,B_n}^0)) \right\}. \quad (3.21)$$

We now study each term separately. For convenience, we use the notation $D_Y(P, \varepsilon) \equiv D_Y^* (\hat{Q}(P)(\varepsilon), \hat{g}(P))$.

The term (3.20) can be written as

$$\begin{aligned}
& E_{B_n} (P_{n,B_n}^1 - P_0) D^* (\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \\
&= E_{B_n} (P_{n,B_n}^1 - P_0) D_Y^* \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{g}(P_{n,B_n}^0) \right) \\
&+ E_{B_n} (P_{n,B_n}^1 - P_0) \left\{ \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - P_W(P_{n,B_n}^1) \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right\} \\
&= E_{B_n} (P_{n,B_n}^1 - P_0) \left\{ D_Y (P_{n,B_n}^0, \varepsilon_n) - D_Y (P_0, \varepsilon_0) \right\} \tag{3.22}
\end{aligned}$$

$$\begin{aligned}
&+ (P_n - P_0) D_Y (P_0, \varepsilon_0) \\
&+ E_{B_n} (P_{n,B_n}^1 - P_0) \left\{ \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(P_0)(\varepsilon_0) \right\} \tag{3.23} \\
&+ (P_n - P_0) \hat{Q}(P_0)(\varepsilon_0).
\end{aligned}$$

It follows from the following lemma that ε_n converges to ε_0 in probability.

Lemma 3.6. *Let ε_n and ε_0 be defined as in Theorem 3.3 and suppose they solve the derivative equations as stated in the theorem. If \hat{g} and \hat{Q} satisfy the conditions 1-5 in Theorem 3.3, then ε_n converges to ε_0 in probability.*

Now consider the following lemmas:

Lemma 3.7. *If the initial estimators \hat{Q} and \hat{g} satisfy the conditions 1-5 in the theorem, then, on a sample split of B_n ,*

$$\sqrt{n}(P_{n,B_n}^1 - P_0) \left\{ D_Y (P_{n,B_n}^0, \varepsilon_n) - D_Y (P_0, \varepsilon_0) \right\} = o_P(1).$$

Lemma 3.8. *If \hat{Q} and \hat{g} satisfy conditions 1-5 of the theorem, then, on a sample split of B_n ,*

$$\sqrt{n}(P_{n,B_n}^1 - P_0) \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(P_0)(\varepsilon_0) \right) = o_P(1).$$

Note that lemmas 3.6, 3.7 and 3.8 follow from lemmas 3.2, 3.4 and 3.5.

Lemmas 3.7 and 3.8 imply that (3.22) and (3.23) are $o_P(1/\sqrt{n})$. We thus have established that (3.20) is given by

$$\begin{aligned}
& E_{B_n} (P_{n,B_n}^1 - P_0) D^* (\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \\
&= (P_n - P_0) \left\{ D_Y^* \left(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0) \right) + \hat{Q}(P_0)(\varepsilon_0) \right\} + o_P(1/\sqrt{n}).
\end{aligned}$$

Combining this result with (3.21), we have proved (3.3):

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= (P_n - P_0) \left\{ D_Y^* \left(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0) \right) + \hat{Q}(P_0)(\varepsilon_0) \right\} \\ &+ E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) (g_0 - \hat{g}(P_{n,B_n}^0)) \right\} \\ &+ o_P(1/\sqrt{n}). \end{aligned}$$

Note that up till this point we have only used convergence of $\hat{Q}(P_n)$ and $\hat{g}(P_n)$ to some limits, but we assumed neither consistency to the true Q_0, g_0 , nor a rate of convergence for these initial estimators to such limits.

Finally, we study the remainder term (3.21):

$$E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) (g_0 - \hat{g}(P_{n,B_n}^0)) \right\}.$$

We consider several cases. Firstly, consider the case, $\hat{g}(P_n) = g_0$. In this case, term (3.21) is exactly 0. Therefore, (3.3) now implies that $\hat{\Psi}(P_n)$ is asymptotically linear with influence curve $D^*(\hat{Q}(P_0)(\varepsilon_0), g_0)$. If in addition, the initial estimator \hat{Q} is consistent for \bar{Q}_0 , i.e. $\hat{Q}(P_0) = \bar{Q}_0$, then

$$\begin{aligned} \varepsilon_0 &\equiv \arg \min_{\varepsilon} P_0 (Y - \hat{Q}(P_0) - \varepsilon H_{\hat{g}(P_0)}^*)^2 \\ &= \arg \min_{\varepsilon} P_0 (Y - Q_0 - \varepsilon H_{g_0}^*)^2 = 0. \end{aligned}$$

This implies that $\hat{Q}(P_0)(\varepsilon_0)$ is simply Q_0 . Consequently, $\hat{\Psi}(P_n)$ is asymptotically linear with influence curve $D^*(Q_0, g_0)$, and is thereby asymptotically efficient.

Let's now consider the case that $\hat{g}(P_0) = g_0$ and $\hat{Q}(P_0) = \bar{Q}_0$. These imply that (3.21) converges to 0. However, for $\hat{\Psi}(P_n)$ to be asymptotically linear, it is necessary that the convergence of this second order term occurs at a \sqrt{n} rate, i.e.

$$\begin{aligned} &E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} (g_0 - \hat{g}(P_{n,B_n}^0)) \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \bar{Q}_0 \right) \right\} \\ &= o_P(1/\sqrt{n}). \end{aligned}$$

Applying Cauchy-Schwartz inequality, it follows that if

$$\begin{aligned} &\sqrt{E_{B_n} P_0 \left(\frac{g_0 - \hat{g}(P_{n,B_n}^0)}{g_0 \hat{g}(P_{n,B_n}^0)} \right)^2} \sqrt{E_{B_n} P_0 \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \bar{Q}_0 \right)^2} \\ &= o_P(1/\sqrt{n}), \end{aligned}$$

then $\hat{\Psi}(P_n)$ will be asymptotically efficient.

Finally, consider the case that $\hat{g}(P_0) = g_0$, but $\hat{Q}(P_0) \neq \bar{Q}_0$. We reconsider the expression (3.21) to account for the limit $\hat{Q}(P_0)(\varepsilon_0)$ of $\hat{Q}(P_{n,B_n}^0)(\varepsilon_n)$ which does not equal \bar{Q}_0 :

$$\begin{aligned} & E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} (g_0 - \hat{g}(P_{n,B_n}^0)) \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \bar{Q}_0 \right) \right\} \\ &= E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} (g_0 - \hat{g}(P_{n,B_n}^0)) \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(P_0)(\varepsilon_0) \right) \right\} \end{aligned} \quad (3.24)$$

$$+ E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} (g_0 - \hat{g}(P_{n,B_n}^0)) \left(\hat{Q}(P_0)(\varepsilon_0) - \bar{Q}_0 \right) \right\}. \quad (3.25)$$

Firstly, we require again that the rate of convergence for the second order term in (3.24) be \sqrt{n} , that is,

$$\begin{aligned} & E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} (g_0 - \hat{g}(P_{n,B_n}^0)) \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(P_0)(\varepsilon_0) \right) \right\} \\ &= o_P(1/\sqrt{n}). \end{aligned}$$

Applying Cauchy-Schwartz inequality, it suffices that

$$\begin{aligned} & \sqrt{E_{B_n} P_0 \left(\frac{g_0 - \hat{g}(P_{n,B_n}^0)}{g_0 \hat{g}(P_{n,B_n}^0)} \right)^2} \sqrt{E_{B_n} P_0 \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(P_0)(\varepsilon_0) \right)^2} \\ &= o_P(1/\sqrt{n}). \end{aligned}$$

For (3.25) to be asymptotically linear, stronger requirements on the performance of \hat{g} are needed in order to address the inconsistency of \hat{Q} . For convenience of notation, we recall that

$$\begin{aligned} & E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) (g_0 - \hat{g}(P_{n,B_n}^0)) \right\} \\ &= E_{B_n} P_0 \left\{ \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{g_0}^* \right) \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) \right\}. \end{aligned}$$

Now, for the given initial estimator \hat{Q} and \hat{g} , let

$$\Phi(P) \equiv P_0 \left\{ H_{\hat{g}(P)}^* \left(\hat{Q}(P_0)(\varepsilon_0) - \bar{Q}_0 \right) \right\}.$$

If \hat{g} is such that $\Phi(P_n) - \Phi(P_0)$ is asymptotically linear (with some influence curve IC'), then (3.25) becomes

$$\begin{aligned} & E_{B_n} P_0 \left\{ \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{g_0}^* \right) \left(\hat{Q}(P_0)(\varepsilon_0) - \bar{Q}_0 \right) \right\} \\ & \equiv E_{B_n} \left(\Phi(P_{n,B_n}^0) - \Phi(P_0) \right) \\ & = E_{B_n} \left(P_{n,B_n}^0 - P_0 \right) IC' + o_P(1/\sqrt{n}) \\ & = (P_n - P_0) IC' + o_P(1/\sqrt{n}). \end{aligned}$$

Therefore, if \hat{g} and \hat{Q} satisfy the convergence speed condition and $\Phi(P_n) - \Phi(P_0)$ asymptotically linear, then it follows from (3.24) and (3.25) that the remainder (3.21) becomes

$$\begin{aligned} & E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} (g_0 - \hat{g}(P_{n,B_n}^0)) \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \bar{Q}_0 \right) \right\} \\ & = (P_n - P_0) IC' + o_P(1/\sqrt{n}). \end{aligned}$$

This completes the proof. \square

Quasi-log-likelihood loss for \bar{Q}

Suppose now that the outcome Y has support in \mathbb{R} and is naturally bounded. After a linear transformation, we may assume without loss of generality that Y has support in $(0, 1)$. Let the loss function be

$$L(\bar{Q})(O) \equiv - (Y \log(\bar{Q}) + (1 - Y) \log(1 - \bar{Q})),$$

and consider the parametric working model through $\bar{Q}(P)$ for any $P \in \mathcal{M}$:

$$\bar{Q}(P)(\varepsilon) = \frac{1}{1 + e^{-\text{logit}(\bar{Q}(P)) - \varepsilon H_{g(P)}^*}}.$$

Then, for the given initial estimators \hat{g} and \hat{Q} , we obtain the following parametric working model:

$$\hat{Q}(P_n)(\varepsilon) = \frac{1}{1 + e^{-\text{logit}(\hat{Q}(P_n)) - \varepsilon H_{\hat{g}(P_n)}^*}}. \quad (3.26)$$

The cross validation selector of ε in (3.26) is defined as

$$\begin{aligned}\varepsilon_n &\equiv \arg \min_{\varepsilon} E_{B_n} P_{n,B_n}^1 L(\hat{Q}(P_{n,B_n}^0)(\varepsilon)) \\ &= \arg \min_{\varepsilon} -E_{B_n} \sum_{i, B_n(i)=1} \left\{ Y_i \log(\hat{Q}(P_{n,B_n}^0)(\varepsilon)(A_i, W_i)) \right. \\ &\quad \left. + (1 - Y) \log(1 - \hat{Q}(P_{n,B_n}^0)(\varepsilon)(A_i, W_i)) \right\}.\end{aligned}$$

At each sample split B_n , we define the TMLE of Q_0 at (P_n, B_n) as

$$\hat{Q}(P_n, B_n)(\varepsilon_n) \equiv \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{Q}_W(P_{n,B_n}^1) \right),$$

where $\hat{Q}_W(P_{n,B_n}^1)$ is the marginal empirical distribution of W in the validation set.

The TMLE of ψ_0 is defined as

$$\hat{\Psi}(P_n) \equiv E_{B_n} \Psi(\hat{Q}(P_n, B_n)(\varepsilon_n) = E_{B_n} \Psi \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n), \hat{Q}_W(P_{n,B_n}^1) \right).$$

Asymptotic results for CV-TMLE under the quasi-log-likelihood loss parallel those for the squared error loss function.

Theorem 3.4. *Consider the setting defined above.*

Suppose that $P_0(|Y| < 1) = 1$.

Let $B_n \in \{0, 1\}^n$ be a random vector indicating a split of $\{1, \dots, n\}$ into a training and validation sample. Suppose B_n is uniformly distributed on a finite support.

Let \hat{Q} and \hat{g} be initial estimators of \bar{Q}_0 and g_0 . In the following, $\hat{Q}(P_0)$ and $\hat{g}(P_0)$ denote limits of these estimators, not necessarily equal to \bar{Q}_0 and g_0 , respectively.

The cross validated TMLE satisfies

$$\begin{aligned}\hat{\Psi}(P_n) - \psi_0 &= E_{B_n} (P_{n,B_n}^1 - P_0) D^* \left(\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0) \right) \\ &+ E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) (g_0 - \hat{g}(P_{n,B_n}^0)) \right\}. \quad (3.2)\end{aligned}$$

Consider the following definition:

$$\varepsilon_0 \equiv \arg \min_{\varepsilon} P_0 L(\hat{Q}(P_0)(\varepsilon)).$$

Suppose that this minimum exists and satisfies the derivative equation

$$0 = P_0 D_Y(P_0, \varepsilon_0),$$

where

$$\begin{aligned} D_Y(P, \varepsilon) &\equiv -\frac{d}{d\varepsilon}L(O, \hat{Q}(P)(\varepsilon)) \\ &= \left(Y - \hat{Q}(P)(\varepsilon)\right) H_{\hat{g}(P)}^* \\ &= D_Y^* \left(\hat{Q}(P)(\varepsilon), \hat{g}(P)\right). \end{aligned}$$

If there are multiple minima, then it is assumed that the argmin is uniquely defined and selects one of these minima.

Suppose that \hat{Q} and \hat{g} satisfy the following conditions:

1. There exists a closed bounded set $K \subset \mathbb{R}^k$ containing ε_0 such that ε_n belongs to K with probability 1;
2. For some $\delta > 0$, $P(1 - \delta > \hat{g}(P_n)(1 | W) > \delta) = 1$;
3. For some $\gamma > 0$, $P(1 - \gamma |\hat{Q}(P_n)(A, W)| < \gamma) = 1$;

4.

$$\int_W (\hat{g}(P_n)(1|w) - \hat{g}(P_0)(1|w))^2 dQ_{W,0}(w) \rightarrow 0 \text{ in probability};$$

5. For $a = 0, 1$,

$$\int_W \left(\hat{Q}(P_n)(a, w) - \hat{Q}(P_0)(a, w)\right)^2 dQ_{W,0}(w) \rightarrow 0 \text{ in probability.}$$

Then,

$$\begin{aligned} \hat{\Psi}(P_n) - \psi_0 &= (P_n - P_0) \left\{ D_Y^* \left(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)\right) + \hat{Q}(P_0)(\varepsilon_0) \right\} \\ &+ E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n)\right) (g_0 - \hat{g}(P_{n,B_n}^0)) \right\} \\ &+ o_P(1/\sqrt{n}). \end{aligned} \quad (3.3)$$

Furthermore, If $\hat{g}(P_n) = g_0$, the TMLE estimator $\hat{\Psi}(P_n)$ is asymptotically linear estimator of ψ_0 :

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) D^* \left(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0)\right) + o_P(1/\sqrt{n}), \quad (3.4)$$

where $\hat{Q}(P_0)(\varepsilon_0) = (\hat{Q}(P_0)(\varepsilon_0), Q_{W,0})$.

If, in addition to $\hat{g}(P_n) = g_0$, $\hat{Q}(P_0) = \bar{Q}_0$, which implies that $\hat{Q}(P_0)(\epsilon_0) = \bar{Q}_0$, then $\hat{\Psi}(P_n)$ is an asymptotically efficient estimator of ψ_0 :

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0)D^*(Q_0, g_0) + o_P(1/\sqrt{n}). \quad (3.5)$$

More generally, if the limits satisfy $\hat{g}(P_0) = g_0$ and $\hat{Q}(P_0) = \bar{Q}_0$, and if the convergence satisfies

$$\begin{aligned} & \sqrt{E_{B_n}P_0 \left(\frac{g_0 - \hat{g}(P_{n,B_n}^0)}{g_0 \hat{g}(P_{n,B_n}^0)} \right)^2} \sqrt{E_{B_n}P_0 \left(\hat{Q}(P_{n,B_n}^0)(\epsilon_n) - \bar{Q}_0 \right)^2} \\ &= o_P(1/\sqrt{n}), \quad (3.16) \end{aligned}$$

then $\hat{\Psi}(P_n)$ is an asymptotically efficient estimator of ψ_0 :

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0)D^*(Q_0, g_0) + o_P(1/\sqrt{n}).$$

Consider now the case that $\hat{g}(P_0) = g_0$, but $\hat{Q}(P_0) \neq \bar{Q}_0$. If the convergence satisfies

$$\begin{aligned} & \sqrt{E_{B_n}P_0 \left(\frac{g_0 - \hat{g}(P_{n,B_n}^0)}{g_0 \hat{g}(P_{n,B_n}^0)} \right)^2} \sqrt{E_{B_n}P_0 \left(\hat{Q}(P_{n,B_n}^0)(\epsilon_n) - \hat{Q}(P_0)(\epsilon_0) \right)^2} \\ &= o_P(1/\sqrt{n}), \quad (3.17) \end{aligned}$$

and $P_0 \left\{ H_{\hat{g}(P_n)}^* \left(\hat{Q}(P_0)(\epsilon_0) - \bar{Q}_0 \right) \right\}$ is an asymptotically linear estimator of the function $P_0 \left\{ H_{\hat{g}(P_0)}^* \left(\hat{Q}(P_0)(\epsilon_0) - \bar{Q}_0 \right) \right\}$ with influence curve IC' , then $\hat{\Psi}(P_n)$ is an asymptotically linear estimator of ψ_0 :

$$\hat{\Psi}(P_n) - \psi_0 = (P_n - P_0) \left\{ D^*(\hat{Q}(P_0)(\epsilon_0), g_0) + IC' \right\} + o_P(1/\sqrt{n}).$$

The proof of this theorem follows the same steps as that of Theorem 3.3. The two only differ in the proofs of some of the auxiliary lemmas. We state the following useful results.

For convenience, we adopt the notation $C_P \equiv \frac{1 - \hat{Q}(P)}{\hat{Q}(P)}$.

Lemma 3.9. Suppose \hat{Q} is such that for some $1 > \gamma > 0$, $P(1 - \gamma > \hat{Q}(P_n)(A, W) > \gamma) = 1$. If for $a = 0, 1$, \hat{g} satisfy

$$\int_W \left(\hat{Q}(a, w) - \hat{Q}(P_0)(a, w) \right)^2 dQ_{W,0}(w) \xrightarrow{P} 0,$$

then

$$P_0(C_{P_n} - C_{P_0})^4 \xrightarrow{P} 0. \quad (3.27)$$

Lemma 3.10. *Suppose \hat{g} and \hat{Q} satisfy the conditions 2-5 in Theorem 3.4. Then, on each split of B_n , for any $r \geq 1$,*

1. $EP_0 \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right)^r \rightarrow 0$;
2. $EP_0 \left(C_{P_{n,B_n}^0} - C_{P_0} \right)^r \rightarrow 0$.

Proof of Theorem 3.4.

The identity in (3.2) is a result of the properties of $\Psi(P)$, its canonical gradient, the definition of ε_n and the one-step convergence of $\hat{Q}(P)(\varepsilon_n)$. Therefore, identical arguments as in the proof of Theorem 3.3 yield (3.2).

We may express the first summand of (3.2) as

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) D^* \left(\hat{Q}(P_{n,B_n})(\varepsilon_n), \hat{g}(P_{n,B_n}^0) \right) \\ &= E_{B_n} (P_{n,B_n}^1 - P_0) \left\{ D_Y(P_{n,B_n}^0, \varepsilon_n) - D_Y(P_0, \varepsilon_0) \right\} \end{aligned} \quad (3.28)$$

$$\begin{aligned} &+ (P_n - P_0) D_Y(P_0, \varepsilon_0) \\ &+ E_{B_n} (P_{n,B_n}^1 - P_0) \left\{ \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(P_0)(\varepsilon_0) \right\} \\ &+ (P_n - P_0) \hat{Q}(P_0)(\varepsilon_0). \end{aligned} \quad (3.29)$$

It follows from the following lemma that ε_n converges to ε_0 in probability.

Lemma 3.11. *Let ε_n and ε_0 be defined as in Theorem 3.4 and suppose they solve the derivative equations as stated in the theorem. If \hat{g} and \hat{Q} satisfy the conditions 1-5 in Theorem 3.4, then ε_n converges to ε_0 in probability.*

The following lemmas 3.12 and 3.13 now prove that (3.28) and (3.29) are $o_P(1/\sqrt{n})$.

Lemma 3.12. *If the initial estimators \hat{Q} and \hat{g} satisfy conditions 1-5 in the theorem, then, on a sample split of B_n ,*

$$\sqrt{n}(P_{n,B_n}^1 - P_0) \left\{ D_Y(P_{n,B_n}^0, \varepsilon_n) - D_Y(P_0, \varepsilon_0) \right\} = o_P(1).$$

Lemma 3.13. *If \hat{Q} and \hat{g} satisfy conditions 1-5 of the theorem, then, on each sample split,*

$$\sqrt{n}(P_{n,B_n}^1 - P_0) \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(P_0)(\varepsilon_0) \right) = o_P(1).$$

Lemmas 3.7 and 3.8 imply that (3.22) and (3.23) are $o_P(1/\sqrt{n})$.

We thus have established that

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) D^* (\hat{Q}(P_n, B_n)(\varepsilon_n), \hat{g}(P_{n,B_n}^0)) \\ &= (P_n - P_0) \left\{ D_Y^* \left(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0) \right) + \hat{Q}(P_0)(\varepsilon_0) \right\} + o_P(1/\sqrt{n}). \end{aligned}$$

Combining this result with the (3.2), we have (3.3).

Finally, we study the remainder term:

$$E_{B_n} P_0 \left\{ \frac{(-1)^{1+A}}{g_0 \hat{g}(P_{n,B_n}^0)} \left(\bar{Q}_0 - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) (g_0 - \hat{g}(P_{n,B_n}^0)) \right\}.$$

Firstly note that if the initial estimator \hat{Q} is consistent for \bar{Q}_0 , i.e. $\hat{Q}(P_0) = \bar{Q}_0$, then $\varepsilon = 0$ is a solution to the derivative equation $P_0 D_Y(\varepsilon, P_0) = 0$. On the other hand, we have seen in the proof of lemma 3.11 that the derivative function is monotonic in ε . Hence, we have $\varepsilon_0 = 0$ and $\hat{Q}(P_0)(\varepsilon_0)$ is simply Q_0 . Now, identical arguments in the proof of Theorem 3.3 complete the proof. \square .

Discussion of conditions of Theorems 3.3 and 3.4.

Under no conditions, we determined an exact identity (3.2), which shows that the analysis of the CV-TMLE involves a cross-validated empirical process term applied to the efficient influence curve, and a second order remainder term. Our second result (3.3) replaces the cross-validated empirical process term by an empirical mean of mean zero random variables $D^*(\hat{Q}(P_0)(\varepsilon_0), \hat{g}(P_0))$ plus a negligible $o_P(1/\sqrt{n})$ -term. That is, under essentially no conditions beyond the positivity assumption, the CV-TMLE minus the true ψ_0 , behaves as an empirical mean of mean zero i.i.d. random variables (which thus converges to a normal distribution, by CLT), plus a specified second order remainder term.

The second order remainder term predicts immediately that to make it negligible we will need that the product of the rates of convergence for $\hat{Q}(P_n)$ and $\hat{g}(P_n)$ to their targets \bar{Q}_0 and g_0 is $o(1/\sqrt{n})$. As mentioned before, in an RCT g_0 is known, so that one might set $\hat{g}(P_n) = g_0$, in which case the second order remainder term is exactly equal to zero, giving us the asymptotic linearity (3.4) of the CV-TMLE under no other conditions than the positivity assumption and convergence of $\hat{Q}(P_n)$ to some fixed function. This teaches us in particular that in an RCT in which we use a consistent estimator \hat{Q} the CV-TMLE is asymptotically efficient, as stated in (3.5). That is, in an RCT, this theorem teaches us that CV-TMLE with adaptive estimation of \bar{Q}_0 is the way to go.

Let's now consider a study in which g_0 is not known, but one has available a correctly specified parametric model: for example, one knows that A is only a function of a discrete

variable, and one uses a saturated model. If the initial estimator \hat{Q} is consistent for \bar{Q}_0 , then the rate condition (3.16) holds, so that it follows that the CV-TMLE is asymptotically efficient. That is, in this scenario there is only benefit in using an adaptive estimator of \bar{Q}_0 . If, by chance, the estimator \hat{Q} is actually inconsistent for \bar{Q}_0 , then the rate condition (3.17) still holds, and the asymptotic linearity condition on \hat{g} will also hold under minimal conditions, so that we still have that the CV-TMLE is asymptotically linear.

Finally, let's consider a case in which the assumed model for g_0 is a large semiparametric model. To have a chance of being consistent for g_0 , one will need to utilize adaptive estimation to estimate g_0 such as a maximum likelihood based super learner respecting the semiparametric model. There are now two scenarios possible. Firstly, suppose that \hat{Q} converges to \bar{Q}_0 fast enough so that (3.16) holds. Then the CV-TMLE is asymptotically efficient. If, on the other hand, \hat{Q} converges fast enough to a misspecified \bar{Q} so that (3.17) holds, then another condition is required. Namely, we now need that \hat{g} is such that the smooth functional $\Phi_{P_0}(\hat{g})$, indexed by P_0 , is an asymptotically linear estimator of its limit $\Phi_{P_0}(g_0)$. This smooth functional can be represented as $\Phi_{P_0}(g) = P_0 H_g^*(\bar{Q}^* - Y)$, where $\bar{Q}^* = \hat{Q}(P_0)(\varepsilon_0)$. A data adaptive estimator \hat{g} of g_0 , only tailored to fit g_0 as a whole, may be too biased for this smooth functional (the whole motivation of TMLE!). Therefore, we suggest that the estimator \hat{g} should be targeted towards this smooth functional. That is, one might want to work out a TMLE \hat{g}^* that aims to target this parameter $\Phi_{P_0}(g_0)$. We leave this for future research.

3.5 The iterative targeted MLE using V-fold sample splitting.

For a given cross-validation scheme $B_n \in \{0, 1\}^n$, we defined

$$\varepsilon_n^0 = \hat{\varepsilon}(P_n) = \arg \min_{\varepsilon} E_{B_n} P_{n, B_n}^1 L(\hat{Q}(P_{n, B_n}^0)(\varepsilon)).$$

This now yields an update $\hat{Q}(P_{n, B_n}^0)(\varepsilon_n^0)$ of $\hat{Q}(P_{n, B_n}^0)$ for each split of B_n . One could now iterate this updating process of the training sample specific estimators: define $\hat{Q}^1(P_{n, B_n}^0) = \hat{Q}(P_{n, B_n}^0)(\varepsilon_n^0)$,

$$\varepsilon_n^1 = \arg \min_{\varepsilon} E_{B_n} P_{n, B_n}^1 L(\hat{Q}^1(P_{n, B_n}^0)(\varepsilon)),$$

resulting in another update $\hat{Q}^1(P_{n, B_n}^0)(\varepsilon_n^1)$ for each B_n . This process is iterated till $\varepsilon_n^K = 0$ (or close enough to zero). We denote the k -step estimator $\hat{Q}^{k-1}(P)(\varepsilon_n^{k-1})$ as $\hat{Q}(P)(\vec{\varepsilon}_n^k)$ to remind us that it is a function of the initial estimators \hat{Q} , \hat{g} and the fluctuation vector

$\vec{\varepsilon}_n^k \equiv (\varepsilon_n^0, \dots, \varepsilon_n^{k-1})$. We denote the k -step TMLE of ψ_0 as

$$\hat{\Psi}^k(P_n) \equiv E_{B_n} \Psi(\hat{Q}(P_{n,B_n}^0)(\vec{\varepsilon}_n^k)).$$

The final update will be denoted with $\hat{Q}(P_{n,B_n}^0)(\vec{\varepsilon}_n^*)$ for each split B_n . The targeted MLE is now defined as $\hat{\Psi}^*(P_n) = E_{B_n} \Psi(\hat{Q}(P_{n,B_n}^0)(\vec{\varepsilon}_n^*))$. We assume that, due to the derivative condition, $\frac{d}{d\varepsilon} L(\hat{Q}(P_n)(\varepsilon)) \Big|_{\varepsilon=0} = D^*(\hat{Q}(P_n), \hat{g}(P_n))$, we have

$$0 = E_{B_n} P_{n,B_n}^1 D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\varepsilon}_n^*), \hat{g}(P_{n,B_n}^0)).$$

We note that $\hat{Q}(P)(\vec{\varepsilon}_n^*)$ is itself dependent on the data through the iterative sequence of selected ε 's: $\varepsilon_n^0, \dots, \varepsilon_n^K$.

We are now ready to present the asymptotics of the k -step cross validated TMLE.

Theorem 3.5. *Let $\hat{Q}(P_n)$, $\hat{g}(P_n)$ be initial estimators of Q_0 , g_0 , respectively, and we will denote their limits with $\hat{Q}(P_0)$ and $\hat{g}(P_0)$, which are not necessarily Q_0 and g_0 , respectively. **Uniformly bounded loss function:** We assume that $\{\hat{Q}(P_n)(\varepsilon) : \varepsilon\} \in \mathcal{Q}$ with probability 1, and that the loss function $L(Q)$ for Q_0 is uniformly bounded in $Q \in \mathcal{Q}$, and over a support of $O \sim P_0$:*

$$M_1 = \sup_Q \sup_O |L(Q)(O)| < \infty.$$

Let $B_n \in \{0, 1\}^n$ be a random vector indicating a split of $\{1, \dots, n\}$ into a training and validation sample. Suppose B_n is uniformly distributed over a finite support.

Suppose there exists $k_n = \hat{k}(P_n) > 0$ such that $P(\hat{k}(P_n) \leq k_0) \rightarrow 1$ for some $k_0 \equiv k(P_0)$ and

$$E_{B_n} P_{n,B_n}^1 D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\varepsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0)) = o_P(1/\sqrt{n}). \quad (3.30)$$

Consider a k_0 -dimensional random vector $\vec{\varepsilon}_n^{k_0} \equiv (\vec{\varepsilon}_n^{k_n}, a_0, \dots, a_0)$, where a_0 is a constant that depends on the choice of the parametric working model such that $\hat{Q}(P)(\vec{\varepsilon}_n^{k_0}) = \hat{Q}(P)(\vec{\varepsilon}_n^{k_n})$. (e.g. $a_0 = 0$ in most cases) Note that $\vec{\varepsilon}_n^{k_n}$ is a projection of $\vec{\varepsilon}_n^{k_0}$ onto its first k_n coordinates.

If parameter $P \rightarrow \Psi(Q(P))$ satisfies

A1:

$$\Psi(Q(P)) - \Psi(Q_0) = -P_0 D^*(Q(P), g_0) + O_P(\|\Psi(Q(P)) - \Psi(Q_0)\|^2).$$

Then

$$\begin{aligned}
\hat{\Psi}^{k_n}(P_n) - \psi_0 &= E_{B_n}(P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0)) \\
&+ E_{B_n} P_0 \left\{ D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), g_0) \right\} \\
&- E_{B_n} P_0 \left\{ D^*(Q_0, \hat{g}(P_{n,B_n}^0)) - D^*(Q_0, g_0) \right\} \\
&+ o_P(1/\sqrt{n}) + O_P(\|\hat{\Psi}^{k_n}(P_n) - \psi_0\|^2).
\end{aligned} \tag{3.31}$$

Let $\vec{\epsilon}_0^{k_0}$ denote the limit of $\vec{\epsilon}_n^{k_0}$ as $n \rightarrow \infty$, that is, $\|\vec{\epsilon}_n^{k_0} - \vec{\epsilon}_0^{k_0}\| \xrightarrow{P} 0$. Suppose the following assumption also holds

A2: (Given $\|\vec{\epsilon}_n^{k_0} - \vec{\epsilon}_0^{k_0}\| \xrightarrow{P} 0$.)

Define the class of functions

$$\mathcal{F}(P_{n,v}^0) \equiv \{O \rightarrow D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), \hat{g}(P_0)) : \vec{\epsilon}\},$$

where the set over which $\vec{\epsilon}$ varies is chosen so that it is a subset of \mathbb{R}^{k_0} and contains $\vec{\epsilon}_n^{k_0}$ with probability tending to 1. In addition, for a deterministic sequence δ_n converging to zero as $n \rightarrow \infty$, we also define the sequence of sub-classes

$$\mathcal{F}_{\delta_n}(P_{n,B_n}^0) \equiv \left\{ f_{\vec{\epsilon}} \in \mathcal{F}(P_{n,B_n}^0) : \|\vec{\epsilon} - \vec{\epsilon}_0^{k_0}\| < \delta_n \right\}.$$

Assume that for deterministic sequence δ_n converging to 0, we have

$$E \text{Entro}(\mathcal{F}_{\delta_n}(P_{n,B_n}^0)) \sqrt{P_0 F^2(\delta_n, P_{n,B_n}^0)} \rightarrow 0 \text{ as } n \rightarrow \infty,$$

where $F(\delta_n, P_{n,B_n}^0)$ is the envelope of $\mathcal{F}_{\delta_n}(P_{n,B_n}^0)$.

Then we can write $\hat{\Psi}^{k_n}(P_n) - \psi_0$ as:

$$\begin{aligned}
\hat{\Psi}^{k_n}(P_n) - \psi_0 &= (P_n - P_0) D^*\left(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), \hat{g}(P_0)\right) + o_P(1/\sqrt{n}) \\
&+ E_{B_n} P_0 \left\{ D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), g_0) \right\} \\
&- E_{B_n} P_0 \left\{ D^*(Q_0, \hat{g}(P_{n,B_n}^0)) - D^*(Q_0, g_0) \right\} \\
&+ O_P(\|\hat{\Psi}^{k_n}(P_n) - \psi_0\|^2).
\end{aligned} \tag{3.32}$$

Furthermore, suppose $\hat{g}(P_n) = g_0$. Then

$$\hat{\Psi}^{k_n}(P_n) - \psi_0 = (P_n - P_0) D^*\left(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), g_0\right) + o_P(1/\sqrt{n}).$$

If, in addition to $\hat{g}(P_n) = g_0$, we also have $\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}) = Q_0$, then $\hat{\Psi}^{k_n}(P_n)$ is in fact asymptotically efficient

$$\hat{\Psi}^{k_n}(P_n) - \psi_0 = (P_n - P_0)D^*(Q_0, g_0) + o_P(1/\sqrt{n}).$$

More generally, suppose $\hat{g}(P_0) = g_0$. Let \tilde{Q} denote the limit of $\hat{Q}(P_n)(\vec{\epsilon}_n^{k_n})$ which is not necessarily Q_0 . Assume in addition

A3:

$$\begin{aligned} & E_{B_n} P_0 \left\{ D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), g_0) \right\} \\ & - E_{B_n} P_0 \left\{ D^*(\tilde{Q}, \hat{g}(P_{n,B_n}^0)) - D^*(\tilde{Q}, g_0) \right\} \\ & = o_P(1/\sqrt{n}). \end{aligned}$$

A4: For for some mean zero function $IC'(P_0) \in L_0^2(P_0)$, we have

$$\begin{aligned} & E_{B_n} P_0 \left\{ D^*(\tilde{Q}, \hat{g}(P_{n,B_n}^0)) - D^*(\tilde{Q}, g_0) \right\} \\ & - E_{B_n} P_0 \left\{ D^*(Q_0, \hat{g}(P_{n,B_n}^0)) - D^*(Q_0, g_0) \right\} \\ & = (P_n - P_0)IC'(P_0) + o_P(1/\sqrt{n}). \end{aligned}$$

NOTE: If $\hat{Q}(P_n)(\vec{\epsilon}_n^{k_n})$ converges to Q_0 then A5 is automatically true with $IC' \equiv 0$.

Then $\hat{\Psi}^{k_n}(P_n)$ is asymptotically linear

$$\hat{\Psi}^{k_n}(P_n) - \psi_0 = (P_n - P_0) \left\{ D^*(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), g_0) + IC'(P_0) \right\} + o_P(1/\sqrt{n}).$$

Proof of Theorem 3.5:

From definition of k_n , we have that

$$E_{B_n} P_{n,B_n}^1 D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0)) = o_P(1/\sqrt{n}).$$

The double robustness of D^* guarantees $P_0 D^*(Q_0, g) = 0$ for all g . Combining this result with A1, we readily have (3.31):

$$\hat{\Psi}^{k_n}(P_n) - \psi_0 = E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0)) \quad (3.33)$$

$$+ E_{B_n} P_0 \left\{ D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), g_0) \right\} \quad (3.34)$$

$$- E_{B_n} P_0 \left\{ D^*(Q_0, \hat{g}(P_{n,B_n}^0)) - D^*(Q_0, g_0) \right\} \quad (3.35)$$

$$+ o_P(1/\sqrt{n}) + O_P(\|\hat{\Psi}^{k_n}(P_n) - \psi_0\|^2).$$

We may rewrite (3.33) as

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(\vec{\epsilon}_n^{k_n}) P_{n,B_n}^0, \hat{g}(P_{n,B_n}^0)) \\ &= E_{B_n} (P_{n,B_n}^1 - P_0) \left\{ D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_0}), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), \hat{g}(P_0)) \right\} \\ &+ E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), \hat{g}(P_0)) \end{aligned}$$

An application of A2 and lemma 3.2, combined with the fact that B_n is uniformly distributed over a finite support, we have

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) \left\{ D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_0}), \hat{g}(P_{n,B_n}^0)) - D^*(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), \hat{g}(P_0)) \right\} \\ &= o_P(1/\sqrt{n}). \end{aligned}$$

In other words, the term (3.33) is given by

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0)) \\ &= E_{B_n} (P_{n,B_n}^1 - P_0) D^*(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), \hat{g}(P_0)) + o_P(1/\sqrt{n}). \end{aligned}$$

This result and the established equality in (3.31) now prove (3.32).

Now, if $\hat{g}(P_n) = g_0$, then the (3.34) and (3.35) are exactly 0. Consequently, (3.32) becomes

$$\begin{aligned} \hat{\Psi}^{k_n}(P_n) - \psi_0 &= (P_n - P_0) D^*(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), g_0) \\ &+ o_P(1/\sqrt{n}) + O_P(\|\hat{\Psi}^{k_n}(P_n) - \psi_0\|^2). \end{aligned}$$

However, note that taking $\|\cdot\|$ on both sides of the equality above yields $\|\hat{\Psi}^{k_n}(P_n) - \psi_0\| = o_P(1/\sqrt{n})$. We thereby have asymptotically linearity of $\hat{\Psi}^{k_n}(P_n)$:

$$\hat{\Psi}^{k_n}(P_n) - \psi_0 = (P_n - P_0) D^*(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), g_0) + o_P(1/\sqrt{n}).$$

If, in addition, $\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}) = Q_0$, then the influence curve is indeed the efficient influence curve $D^*(Q_0, g_0)$.

Next we consider a more general case where $\hat{g}(P_0) = g_0$. Let \tilde{Q} be the limit of $\hat{Q}(P_n)(\vec{\epsilon}_n^{k_n})$. It is not necessarily the case that $\tilde{Q} = Q_0$. We now rewrite the established

equality (3.32) to account for \tilde{Q} :

$$\begin{aligned} \hat{\Psi}^{k_n}(P_n) - \psi_0 &= (P_n - P_0) D^* \left(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), \hat{g}(P_0) \right) + o_P(1/\sqrt{n}) \\ &+ E_{B_n} P_0 \left\{ D^* \left(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0) \right) - D^* \left(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), g_0 \right) \right\} \\ &- E_{B_n} P_0 \left\{ D^* \left(\tilde{Q}, \hat{g}(P_{n,B_n}^0) \right) - D^* \left(\tilde{Q}, g_0 \right) \right\} \\ &+ E_{B_n} P_0 \left\{ D^* \left(\tilde{Q}, \hat{g}(P_{n,B_n}^0) \right) - D^* \left(\tilde{Q}, g_0 \right) \right\} \\ &- E_{B_n} P_0 \left\{ D^* \left(Q_0, \hat{g}(P_{n,B_n}^0) \right) - D^* \left(Q_0, g_0 \right) \right\} \\ &+ O_P(\| \hat{\Psi}^{k_n}(P_n) - \psi_0 \|^2). \end{aligned}$$

From A3, the term

$$\begin{aligned} &E_{B_n} P_0 \left\{ D^* \left(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), \hat{g}(P_{n,B_n}^0) \right) - D^* \left(\hat{Q}(P_{n,B_n}^0)(\vec{\epsilon}_n^{k_n}), g_0 \right) \right\} \\ &- E_{B_n} P_0 \left\{ D^* \left(\tilde{Q}, \hat{g}(P_{n,B_n}^0) \right) - D^* \left(\tilde{Q}, g_0 \right) \right\} \\ &= o_P(1/\sqrt{n}). \end{aligned}$$

From A4, the term

$$\begin{aligned} &E_{B_n} P_0 \left\{ D^* \left(\tilde{Q}, \hat{g}(P_{n,B_n}^0) \right) - D^* \left(\tilde{Q}, g_0 \right) \right\} \\ &- E_{B_n} P_0 \left\{ D^* \left(Q_0, \hat{g}(P_{n,B_n}^0) \right) - D^* \left(Q_0, g_0 \right) \right\} \\ &= (P_n - P_0) IC'(P_0) + o_P(1/\sqrt{n}). \end{aligned}$$

Therefore (3.3) becomes

$$\begin{aligned} \hat{\Psi}^{k_n}(P_n) - \psi_0 &= (P_n - P_0) \left\{ D^* \left(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), \hat{g}(P_0) \right) + IC'(P_0) \right\} + o_P(1/\sqrt{n}) \\ &+ O_P(\| \hat{\Psi}^{k_n}(P_n) - \psi_0 \|^2). \end{aligned}$$

Taking $\| \cdot \|$ on both sides again yields $\| \hat{\Psi}^{k_n}(P_n) - \psi_0 \| = o_P(1/\sqrt{n})$. We thereby have the desired result

$$\hat{\Psi}^{k_n}(P_n) - \psi_0 = (P_n - P_0) \left\{ D^* \left(\hat{Q}(P_0)(\vec{\epsilon}_0^{k_0}), \hat{g}(P_0) \right) + IC'(P_0) \right\} + o_P(1/\sqrt{n}).$$

□

3.6 Concluding remarks.

We presented a TMLE that allows to learn the truth ψ_0 , while also providing statistical inference based on an CLT, under an as large statistical model as possible. For that purpose,

the combination of adaptive estimation (super learning), targeted maximum likelihood estimation, and cross-validated selection of the fluctuation parameter in the TMLE, are essential tools to achieve this goal.

In future work we wish to investigate the extension of CV-TMLE to collaborative targeted maximum likelihood estimation, as in van der Laan and Gruber (2010), and the incorporation of targeted estimators of g_0 to enhance the asymptotic linearity of the CV-TMLE of ψ_0 for the case that the initial estimator of Q_0 is inconsistent.

3.7 Appendix

Proof of lemma 3.2: Let $G_{n,B_n}^1 = \sqrt{n}(P_{n,B_n}^1 - P_0)$. For any $\delta > 0$.

$$\begin{aligned}
P(|G_{n,B_n}^1 f_{\varepsilon_n}(P_{n,B_n}^0)| > \delta) &= EP \left(|G_{n,B_n}^1 f_{\varepsilon_n}(P_{n,B_n}^0)| > \delta \middle| P_{n,B_n}^0 \right) \\
&= EP \left(|G_{n,B_n}^1 f_{\varepsilon_n}(P_{n,B_n}^0) I(\|\varepsilon_n - \varepsilon_0\| < \delta_n)| > \delta \middle| P_{n,B_n}^0 \right) \\
&\quad + EP \left(|G_{n,B_n}^1 f_{\varepsilon_n}(P_{n,B_n}^0) I(\|\varepsilon_n - \varepsilon_0\| \geq \delta_n)| > \delta \middle| P_{n,B_n}^0 \right) \\
&\leq EP \left(\sup_{f \in \mathcal{F}_{\delta_n}(P_{n,B_n}^0)} |G_{n,B_n}^1 f| > \delta \middle| P_{n,B_n}^0 \right) \\
&\quad + EP \left(\|\varepsilon_n - \varepsilon_0\| \geq \delta_n \middle| P_{n,B_n}^0 \right) \\
&= EP \left(\sup_{f \in \mathcal{F}_{\delta_n}(P_{n,B_n}^0)} |G_{n,B_n}^1 f| > \delta \middle| P_{n,B_n}^0 \right) \\
&\quad + P(\|\varepsilon_n - \varepsilon_0\| \geq \delta_n).
\end{aligned}$$

By our assumption, $P(\|\varepsilon_n - \varepsilon_0\| \geq \delta_n) \rightarrow 0$. On the other hand, by Chebysev inequality, lemma 3.1 and Cauchy-Schwartz inequality

$$\begin{aligned}
&EP \left(\sup_{f \in \mathcal{F}_{\delta_n}(P_{n,B_n}^0)} |G_{n,B_n}^1 f| > \delta \middle| P_{n,B_n}^0 \right) \\
&\leq \frac{1}{\delta} EE \left(\sup_{f \in \mathcal{F}_{\delta_n}(P_{n,B_n}^0)} |G_{n,B_n}^1 f| \right) \\
&\leq \frac{1}{\delta} EEntro(\mathcal{F}_{\delta_n}(P_{n,B_n}^0)) \sqrt{P_0 F(\delta_n, P_{n,B_n}^0)^2}.
\end{aligned}$$

Therefore, $G_{n,B_n}^1 f_{\varepsilon_n}(P_{n,B_n}^0) \xrightarrow{P} 0$ by our assumption. \square

Proof of lemma 3.3:

By our assumption, $P(|X_n - X| < 2\eta) = 1$. Then for any $\delta > 0$,

$$\begin{aligned} E|X_n - X|^r &= E\{|X_n - X|^r I_{|X_n - X| \leq \delta}\} + E\{|X_n - X|^r I_{|X_n - X| > \delta}\} \\ &\leq \delta^r P(|X_n - X| \leq \delta) + (2A)^r P(|X_n - X| > \delta) \\ &= \delta^r \cdot 1 + ((2\eta)^r - \delta^r) P(|X_n - X| > \delta). \end{aligned}$$

We assumed that $P(|X_n - X| > \delta) \rightarrow 0$. Hence the last equality converges to δ^r . This holds for all $\delta > 0$. Thus we must have $E|X_n - X|^r \rightarrow 0$. \square

Proof of lemma 3.4:

First note that

$$H_{\hat{g}(P_n)}^*(A, W) - H_{\hat{g}(P_0)}^*(A, W) = (-1)^{A+1} \frac{(\hat{g}(P_0)(A|W) - \hat{g}(P_{n,B_n}^0)(A|W))}{\hat{g}(P_0)(A|W)\hat{g}(P_{n,B_n}^0)(A|W)}. \quad (3.36)$$

The expression $P_0 \left(H_{\hat{g}(P_n)}^* - H_{\hat{g}(P_0)}^* \right)^4$ can be expanded into

$$P_0 \left(H_{\hat{g}(P_n)}^* - H_{\hat{g}(P_0)}^* \right)^4 = \sum_{a=0,1} E_{W,0} \left(H_{\hat{g}(P_n)}^*(a, W) - H_{\hat{g}(P_0)}^*(a, W) \right)^4 g_0(a|W).$$

Applying Cauchy-Schwartz and (3.36), each summand can be bounded as follows:

$$\begin{aligned} &E_{W,0} \left(H_{\hat{g}(P_n)}^*(a, W) - H_{\hat{g}(P_0)}^*(a, W) \right)^4 g_0(a|W) \\ &\leq \sqrt{E_{W,0} \left(\frac{g_0(a|W)}{(\hat{g}(P_0)\hat{g}(P_n)(a|W))^4} \right)^2} \sqrt{E_{W,0} \{ \hat{g}(P_0)(a|W) - \hat{g}(P_n)(a|W) \}^8} \\ &\leq \sqrt{E_{W,0} \left(\frac{g_0(a|W)}{(\hat{g}(P_0)\hat{g}(P_n)(a|W))^4} \right)^2} \sqrt{E_{W,0} \{ \hat{g}(P_0)(a|W) - \hat{g}(P_n)(a|W) \}^2}. \end{aligned}$$

Since $E_{W,0} \left(\frac{g_0(a|W)}{(\hat{g}(P_0)\hat{g}(P_n)(a|W))^4} \right)^2$ is bounded and, by assumption,

$$E_{W,0} (\hat{g}(P_0)(a|W) - \hat{g}(P_n)(a|W))^2 \xrightarrow{P} 0$$

this inequality implies that

$$P_0 \left(H_{\hat{g}(P_n)}^* - H_{\hat{g}(P_0)}^* \right)^4 \xrightarrow{P} 0. \quad (3.37)$$

To prove $P_0 \left(H_{\hat{g}(P_n)}^* - H_{\hat{g}(P_0)}^* \right)^2 \xrightarrow{P} 0$, we use a simple application of Cauchy-Schwartz inequality and (3.37). Similarly for $P_0 \left(H_{\hat{g}(P_n)}^* - H_{\hat{g}(P_0)}^* \right) \xrightarrow{P} 0$.

Finally, to verify $P_0 \left((H_{\hat{g}(P_n)}^*)^2 - (H_{\hat{g}(P_0)}^*)^2 \right) \xrightarrow{P} 0$, we first bound the expectation using Cauchy-Schwartz inequality:

$$\begin{aligned} & \left| P_0 \left((H_{\hat{g}(P_n)}^*)^2 - (H_{\hat{g}(P_0)}^*)^2 \right) \right| \\ & \leq \sqrt{P_0 \left(H_{\hat{g}(P_n)}^* - H_{\hat{g}(P_0)}^* \right)^2} \sqrt{P_0 \left(H_{\hat{g}(P_n)}^* + H_{\hat{g}(P_0)}^* \right)^2}. \end{aligned}$$

By assumption, $P_0 \left(H_{\hat{g}(P_n)}^* + H_{\hat{g}(P_0)}^* \right)^2$ is bounded; on the other hand, we established that that $P_0 \left(H_{\hat{g}(P_n)}^* - H_{\hat{g}(P_0)}^* \right)^2$ converges to 0 in probability. Thus, the above inequality implies that $P_0 \left((H_{\hat{g}(P_n)}^*)^2 - (H_{\hat{g}(P_0)}^*)^2 \right)$ converges to 0 in probability. \square

Proof of lemma 3.5:

1. Firstly, note that

$$\begin{aligned} & E_{B_n} P_0 \left(\hat{Q}(P_{n,B_n}^0) H_{\hat{g}(P_{n,B_n}^0)}^* - \hat{Q}(P_0) H_{\hat{g}(P_0)}^* \right) \\ & = E_{B_n} P_0 \hat{Q}(P_{n,B_n}^0) \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right) \\ & \quad + E_{B_n} P_0 H_{\hat{g}(P_0)}^* \left(\hat{Q}(P_{n,B_n}^0) - \hat{Q}(P_0) \right). \end{aligned}$$

By our assumptions, $P_0 \hat{Q}(P_n)^2$ is bounded with probability 1. Hence, it follows from Cauchy-Schwartz inequality and lemma 3.4 that the first summand converges to 0 in probability. On the other hand, we assumed that $P_0 (H_{\hat{g}(P_0)}^*)^2$ is bounded and $P_0 \left(\hat{Q}(P_n) - \hat{Q}(P_0) \right)^2 \xrightarrow{P} 0$. Therefore, it follows from an application of Cauchy-Schwartz inequality that the second summand also converge to 0. From these two results it follows that

$$E_{B_n} P_0 \left\{ \left(\hat{Q}(P_{n,B_n}^0) H_{\hat{g}(P_{n,B_n}^0)}^* - \hat{Q}(P_0) H_{\hat{g}(P_0)}^* \right) \right\} \xrightarrow{P} 0. \quad (3.38)$$

Secondly, note also that our assumptions imply that $E_{B_n} P_0 \hat{Q}(P_{n,B_n}^0) H_{\hat{g}(P_{n,B_n}^0)}^*$ is bounded with probability 1. Hence, by lemma 3.3 we have obtain the desired result.

2. Firstly, note that

$$\begin{aligned} & E_{B_n} P_0 \left((Y - \hat{Q}(P_{n,B_n}^0)) H_{\hat{g}(P_{n,B_n}^0)}^* - (Y - \hat{Q}(P_0)) H_{\hat{g}(P_0)}^* \right) \\ &= E_{B_n} P_0 \left\{ Y \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right) \right\} \\ & \quad - E_{B_n} P_0 \left\{ \left(\hat{Q}(P_{n,B_n}^0) H_{\hat{g}(P_{n,B_n}^0)}^* - \hat{Q}(P_0) H_{\hat{g}(P_0)}^* \right) \right\}. \end{aligned}$$

By our assumption, $P_0 Y^2$ is bounded. Hence, it follows from Cauchy-Schwartz inequality and lemma 3.4 that $E_{B_n} P_0 \left\{ Y \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right) \right\}$ converges to 0 in probability. Combining this result and (3.38), we have

$$E_{B_n} P_0 \left((Y - \hat{Q}(P_{n,B_n}^0)) H_{\hat{g}(P_{n,B_n}^0)}^* - (Y - \hat{Q}(P_0)) H_{\hat{g}(P_0)}^* \right) \xrightarrow{P} 0.$$

On the other hand, by our assumption, $E_{B_n} P_0 (Y - \hat{Q}(P_{n,B_n}^0)) H_{\hat{g}(P_{n,B_n}^0)}^*$ is bounded with probability 1. Hence, an application of lemma 3.3 yields the desired result.

3. By our assumption, $E_{B_n} P_0 (H_{\hat{g}(P_{n,B_n}^0)}^*)^2$ is bounded with probability 1. Hence, by lemma 3.4 and lemma 3.3, we have $E P_0 \left((H_{\hat{g}(P_{n,B_n}^0)}^*)^2 - (H_{\hat{g}(P_0)}^*)^2 \right)^r \rightarrow 0$ for any $r \geq 1$.
4. Similarly, by our assumption, $E_{B_n} P_0 H_{\hat{g}(P_{n,B_n}^0)}^*$ is bounded with probability 1. Hence by lemma 3.4 and lemma 3.3, we have

$$E P_0 \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right)^r \rightarrow 0$$

for any $r \geq 1$.

□

Proof of lemma 3.6:

By our definition of ε_n and the one-step convergence of the fluctuation model,

$$E_{B_n} P_{n,B_n}^1 D_Y(P_{n,B_n}^0, \varepsilon_n) = 0.$$

This implies that

$$-P_0 D_Y(P_0, \varepsilon_n) = E_{B_n} (P_{n,B_n}^1 - P_0) \{D_Y(P_{n,B_n}^0, \varepsilon_n) - D_Y(P_0, \varepsilon_n)\} \quad (3.39)$$

$$+ E_{B_n} P_0 \{D_Y(P_{n,B_n}^0, \varepsilon_n) - D_Y(P_0, \varepsilon_n)\} \quad (3.40)$$

$$+ E_{B_n} (P_{n,B_n}^1 - P_0) D_Y(P_0, \varepsilon_n) \quad (3.41)$$

The term (3.40) can be expanded into

$$\begin{aligned} & E_{B_n} P_0 \{D_Y(P_{n,B_n}^0, \varepsilon_n) - D_Y(P_0, \varepsilon_n)\} \\ &= E_{B_n} P_0 \left\{ Y \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right) \right\} \\ & - E_{B_n} P_0 \left\{ \left(\hat{Q}(P_{n,B_n}^0) H_{\hat{g}(P_{n,B_n}^0)}^* - \hat{Q}(P_0) H_{\hat{g}(P_0)}^* \right) \right\} \\ & - \varepsilon_n E_{B_n} P_0 \left\{ \left((H_{\hat{g}(P_{n,B_n}^0)}^*)^2 - (H_{\hat{g}(P_0)}^*)^2 \right) \right\}. \end{aligned}$$

Note that ε_n is bounded with probability 1. Therefore, applying the arguments in the proof of lemma 3.5 to the corresponding summands, we have that (3.40) converges to 0 in probability.

The term (3.41) can be written as

$$\begin{aligned} & E_{B_n} (P_{n,B_n}^1 - P_0) D_Y(P_0, \varepsilon_n) \\ & \equiv E_{B_n} (P_{n,B_n}^1 - P_0) (Y - \hat{Q}(P_0) - \varepsilon_n H_{\hat{g}(P_0)}^*) H_{\hat{g}(P_0)}^* \\ & = E_{B_n} (P_{n,B_n}^1 - P_0) Y H_{\hat{g}(P_0)}^* - E_{B_n} (P_{n,B_n}^1 - P_0) \hat{Q}(P_0) H_{\hat{g}(P_0)}^* \\ & - \varepsilon_n E_{B_n} (P_{n,B_n}^1 - P_0) (H_{\hat{g}(P_0)}^*)^2. \end{aligned}$$

All the empirical differences in the last equality are asymptotically normal with mean 0, and ε_n is bounded with probability 1. Therefore, we have that (3.41) converges to 0 in probability.

It remains to show that (3.39) converges to 0 in probability. By our assumption, there exists constant $M > 0$ such that $P(|\varepsilon_n| < M) = 1$. Conditional on P_{n,B_n}^0 , consider the class

$$\mathcal{F}(P_{n,B_n}^0) = \{f_\varepsilon(P_{n,B_n}^0) = D_Y(P_{n,B_n}^0, \varepsilon) - D_Y(P_0, \varepsilon) : |\varepsilon| < M\}.$$

Lemma 3.1 implies that

$$\sqrt{n} E \left(\sup_{f \in \mathcal{F}(P_{n,B_n}^0)} |(P_{n,B_n}^1 - P_0) f| \right) \leq \text{Entro}(\mathcal{F}(P_{n,B_n}^0)) \sqrt{P_0 \mathbf{F}(P_{n,B_n}^0)^2},$$

where $\mathbf{F}(P_{n,B_n}^0)$ is an envelope of $\mathcal{F}(P_{n,B_n}^0)$. Therefore, after an application of Chebysev inequality we may write

$$\begin{aligned} & P\left(|(P_{n,B_n}^1 - P_0)f_{\varepsilon_n}(P_{n,B_n}^0)| > \delta\right) \\ & \leq EP\left(\sup_{f \in \mathcal{F}(P_{n,B_n}^0)} |(P_{n,B_n}^1 - P_0)f| > \delta \middle| P_{n,B_n}^0\right) \\ & \leq \frac{1}{\delta} EE\left(\sup_{f \in \mathcal{F}(P_{n,B_n}^0)} |(P_{n,B_n}^1 - P_0)f|\right) \\ & \leq \frac{1}{\sqrt{n}} \frac{1}{\delta} EE\text{Entro}(\mathcal{F}(P_{n,B_n}^0)) \sqrt{P_0 \mathbf{F}(P_{n,B_n}^0)^2}. \end{aligned}$$

Firstly note that f_ε is bounded per our assumptions. Hence $\sqrt{P_0 \mathbf{F}(P_{n,B_n}^0)^2}$ is bounded. On the other hand, the entropy of the class is also bounded. Therefore, we indeed have $P\left(|(P_{n,B_n}^1 - P_0)f(\varepsilon_n)(P_0)| > \delta\right)$ converges to 0 as $n \rightarrow \infty$. Consequently, (3.39) converges to 0 in probability.

Since K is compact, there is a subsequence ε_{nk} such that $\varepsilon_{nk} \xrightarrow{P} \varepsilon^*$ for some $\varepsilon^* \in K$. This implies that for

$$\begin{aligned} g(\varepsilon) & \equiv P_0 D_Y(P_0, \varepsilon) \\ & = P_0 Y H_{\hat{g}(P_0)}^* - P_0 \hat{Q}(P_0) H_{\hat{g}(P_0)}^* - \varepsilon P_0 (H_{\hat{g}(P_0)}^*)^2, \end{aligned}$$

which is continuous over K , we must have $g(\varepsilon_{nk}) \xrightarrow{P} g(\varepsilon^*)$.

We determined in above that $g(\varepsilon_n) \xrightarrow{P} 0$, therefore it follows that $g(\varepsilon^*) = 0$. On the other hand, by definition of ε_0 we have that $g(\varepsilon_0) = 0$. Since $g(\varepsilon)$ is a linear function in ε , it has unique solution at ε_0 , therefore we indeed have $\varepsilon^* = \varepsilon_0$. This implies that all convergent subsequences of ε_n converge to ε_0 in probability. Since K is compact, it now implies that ε_n converge to ε_0 in probability. \square

Proof of lemma 3.7: Conditional on P_{n,B_n}^0 , for a deterministic sequence δ_n converging to 0, consider the class

$$\mathcal{F}_{\delta_n}(P_{n,B_n}^0) \equiv \left\{ D_Y(P_{n,B_n}^0, \varepsilon) - D_Y(P_0, \varepsilon_0) : \|\varepsilon - \varepsilon_0\| < \delta_n \right\}.$$

From lemma 3.6, we know that $\|\varepsilon_n - \varepsilon_0\| \xrightarrow{P} 0$. To obtain the proposed result, we will show that this class satisfies the conditions of lemma 3.2.

For convenience, let $A_{P_{n,B_n}^0}(O) \equiv \left(Y - \hat{Q}(P_{n,B_n}^0)(A, W) \right) H_{\hat{g}(P_{n,B_n}^0)}^*(A, W)$, $H_{P_{n,B_n}^0}(O)^2 \equiv (H_{\hat{g}(P_{n,B_n}^0)}^*)^2$, and $A_{P_0}, H_{P_0}^2$ denote the analogous functions trained at P_0 . Then, we can find

an envelope for this class of functions as follows:

$$\begin{aligned}
& \left| D_Y(P_{n,B_n}^0, \varepsilon) - D_Y(P_0, \varepsilon_0) \right| \equiv \left| (A_{P_{n,B_n}^0} - \varepsilon H_{P_{n,B_n}^0}^2) - (A_{P_0} - \varepsilon_0 H_{P_0}^2) \right| \\
& = \left| (A_{P_{n,B_n}^0} - A_{P_0}) - H_{P_{n,B_n}^0}^2 (\varepsilon - \varepsilon_0) - \varepsilon_0 (H_{P_{n,B_n}^0}^2 - H_{P_0}^2) \right| \\
& \leq |A_{P_{n,B_n}^0} - A_{P_0}| + |H_{P_{n,B_n}^0}^2| |\delta_n + \varepsilon_0| |H_{P_{n,B_n}^0}^2 - H_{P_0}^2| \\
& \equiv \mathbf{F}_n.
\end{aligned}$$

Now, we study the convergence of $EP_0(\mathbf{F}_n)^2$. From the proposed conditions and lemma 3.5, we readily have that:

$$EP_0(A_{P_{n,B_n}^0} - A_{P_0})^2 \rightarrow 0,$$

and

$$EP_0 \left\{ \varepsilon_0^2 (H_{P_{n,B_n}^0}^2 - H_{P_0}^2)^2 \right\} = \varepsilon_0^2 EP_0 (H_{P_{n,B_n}^0}^2 - H_{P_0}^2)^2 \rightarrow 0.$$

On the other hand, the boundedness conditions for \hat{g} imply that $EP_0 H_{P_{n,B_n}^0}^4$ is bounded. Since δ_n converges to 0, this now implies that

$$EP_0 \left\{ (H_{P_{n,B_n}^0}^2)^4 \delta_n^2 \right\} \rightarrow 0.$$

Thus, all the square terms of $EP_0(\mathbf{F}_n)^2$ converge to 0 as $n \rightarrow \infty$. Applying Cauchy-Schwartz inequality and lemma 3.5 in a similar manner to the cross terms of $EP_0(\mathbf{F}_n)^2$ will show that they also converge to 0.

Moreover, this class has bounded entropy since the functions are linear in ε . Therefore, lemma 3.2 implies that we indeed have the desired result:

$$\sqrt{n}(P_{n,B_n}^1 - P_0) \{ D_Y(P_{n,B_n}^0, \varepsilon_n) - D_Y(P_0, \varepsilon_0) \} = o_P(1).$$

□

Proof of lemma 3.8:

This result can be proved in a similar manner as lemma 3.7 by making use lemma 3.2 and the conditions of the theorem. □

Proof of lemma 3.9: Firstly, rewrite

$$C_{P_n} - C_{P_0} \equiv \frac{\hat{Q}(P_0) - \hat{Q}(P_n)}{\hat{Q}(P_n)\hat{Q}(P_0)}.$$

Then

$$\begin{aligned}
P_0(C_{P_n} - C_{P_0})^4 &= P_0 \frac{(\hat{Q}(P_0) - \hat{Q}(P_n))^4}{(\hat{Q}(P_n)\hat{Q}(P_0))^4} \\
&\leq \sqrt{P_0 \frac{1}{(\hat{Q}(P_n)\hat{Q}(P_0))^8}} \sqrt{P_0 (\hat{Q}(P_0) - \hat{Q}(P_n))^8} \\
&\leq \sqrt{P_0 \frac{1}{(\hat{Q}(P_n)\hat{Q}(P_0))^8}} \sqrt{P_0 (\hat{Q}(P_0) - \hat{Q}(P_n))^2},
\end{aligned}$$

where the last inequality follows from the assumption that $\hat{Q}(P)$ are bounded between 0 and 1 with probability 1. This last expression converges to 0 in probability by our assumption. \square .

Proof of lemma 3.10:

1. By our assumption, $E_{B_n} P_0 H_{\hat{g}(P_{n,B_n}^0)}^*$ is bounded with probability 1. Cauchy-Schwartz inequality and lemma 3.4 imply that $E_{B_n} P_0 \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right) \xrightarrow{P} 0$. It now follows from lemma 3.3 that

$$E P_0 \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right)^r \rightarrow 0$$

for any $r \geq 1$.

2. Similarly, By our assumption, $E_{B_n} P_0 C_{P_{n,B_n}^0}$ is bounded with probability 1. An application of Cauchy-Schwartz inequality and lemma 3.9 implies that $E_{B_n} P_0 (C_{P_{n,B_n}^0} - C_{P_0}) \xrightarrow{P} 0$. Hence, lemma 3.3 yields

$$E P_0 \left(C_{P_{n,B_n}^0} - C_{P_0} \right)^r \rightarrow 0$$

for any $r \geq 1$.

\square

Proof of Lemma 3.11: By our definition,

$$E_{B_n} P_{n,B_n}^1 D_Y(P_{n,B_n}^0, \varepsilon_n) = 0.$$

This implies that

$$-P_0 D_Y(P_0, \varepsilon_n) = E_{B_n} (P_{n,B_n}^1 - P_0) \{D_Y(P_{n,B_n}^0, \varepsilon_n) - D_Y(P_0, \varepsilon_n)\} \quad (3.42)$$

$$+ E_{B_n} P_0 \{D_Y(P_{n,B_n}^0, \varepsilon_n) - D_Y(P_0, \varepsilon_n)\} \quad (3.43)$$

$$+ E_{B_n} (P_{n,B_n}^1 - P_0) D_Y(P_0, \varepsilon_n). \quad (3.44)$$

The term (3.43) can be expanded into

$$\begin{aligned} & E_{B_n} P_0 \{D_Y(P_{n,B_n}^0, \varepsilon_n) - D_Y(P_0, \varepsilon_n)\} \\ &= E_{B_n} P_0 \left\{ \left(Y - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) H_{\hat{g}(P_{n,B_n}^0)}^* - \left(Y - \hat{Q}(\varepsilon_n)(P_0) \right) H_{\hat{g}(P_0)}^* \right\} \\ &= E_{B_n} P_0 \left\{ \left(Y - \hat{Q}(P_{n,B_n}^0)(\varepsilon_n) \right) \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right) \right\} \\ &\quad - E_{B_n} P_0 \left\{ H_{\hat{g}(P_0)}^* \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(\varepsilon_n)(P_0) \right) \right\}. \end{aligned}$$

By lemma 3.4, $E_{B_n} P_0 \left(H_{\hat{g}(P_{n,B_n}^0)}^* - H_{\hat{g}(P_0)}^* \right)^2 \xrightarrow{P} 0$. Moreover, Y is bounded by assumption and $\hat{Q}(P)(\varepsilon)$ is bounded by construction. Hence, an application of Cauchy-Schwartz imply that the first summand of the last equality converges to zero in probability. On the other hand the second summand can be bounded by

$$\begin{aligned} & \left| E_{B_n} P_0 \left\{ H_{\hat{g}(P_0)}^* \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(\varepsilon_n)(P_0) \right) \right\} \right| \\ & \leq \sqrt{E_{B_n} P_0 (H_{\hat{g}(P_0)}^*)^2} \sqrt{E_{B_n} P_0 \left(\hat{Q}(P_{n,B_n}^0)(\varepsilon_n) - \hat{Q}(\varepsilon_n)(P_0) \right)^2} \\ & = \sqrt{E_{B_n} P_0 (H_{\hat{g}(P_0)}^*)^2} \sqrt{E_{B_n} P_0 \frac{\left(C_{P_0} e^{-\varepsilon_n H_{\hat{g}(P_0)}^*} - C_{P_{n,B_n}^0} e^{-\varepsilon_n H_{\hat{g}(P_{n,B_n}^0)}^*} \right)^2}{(1 + C_{P_0} e^{-\varepsilon_n H_{\hat{g}(P_0)}^*})^2 (1 + C_{P_{n,B_n}^0} e^{-\varepsilon_n H_{\hat{g}(P_{n,B_n}^0)}^*})^2}} \\ & \leq \sqrt{E_{B_n} P_0 (H_{\hat{g}(P_0)}^*)^2} \sqrt{E_{B_n} P_0 \left(C_{P_0} e^{-\varepsilon_n H_{\hat{g}(P_0)}^*} - C_{P_{n,B_n}^0} e^{-\varepsilon_n H_{\hat{g}(P_{n,B_n}^0)}^*} \right)^2}. \end{aligned}$$

By our assumption $E_{B_n} P_0 (H_{\hat{g}(P_0)}^*)^2$ is bounded. We now wish to show

$$E_{B_n} P_0 \left(C_{P_{n,B_n}^0} e^{-\varepsilon_n H_{\hat{g}(P_{n,B_n}^0)}^*} - C_{P_0} e^{-\varepsilon_n H_{\hat{g}(P_0)}^*} \right)^2 \xrightarrow{P} 0.$$

Let $H_{P_{n,B_n}^0} \equiv H_{\hat{g}(P_{n,B_n}^0)}^*$, $H_{P_0} \equiv H_{\hat{g}(P_0)}^*$. Firstly, note that by property of the exponential function for every (a, w) in the support, there is $Y_{P_{n,B_n}^0}(a, w)$ between $\varepsilon_n H_{P_{n,B_n}^0}(a, w)$ and $\varepsilon_n H_{P_0}(a, w)$ such that

$$\begin{aligned} e^{\varepsilon_n H_{P_{n,B_n}^0}(a, w)} - e^{\varepsilon_n H_{P_0}(a, w)} &= e^{-\varepsilon_n H_{P_0}(a, w)} \varepsilon_n (H_{P_{n,B_n}^0} - H_{P_0})(a, w) \\ &+ \frac{e^{Y_{P_{n,B_n}^0}(a, w)}}{2} \varepsilon_n^2 (H_{P_{n,B_n}^0} - H_{P_0})^2(a, w). \end{aligned}$$

Boundedness of ε_n and H_{P_{n,B_n}^0} implies that Y_{P_{n,B_n}^0} is also bounded with probability 1 over the support. Therefore, we have:

$$\begin{aligned} &E_{B_n} P_0 \left(C_{P_{n,B_n}^0} e^{-\varepsilon_n H_{P_{n,B_n}^0}} - C_{P_0} e^{-\varepsilon_n H_{P_0}} \right)^2 \\ &= E_{B_n} P_0 \left\{ C_{P_{n,B_n}^0} \left(e^{-\varepsilon_n H_{P_{n,B_n}^0}} - e^{-\varepsilon_n H_{P_0}} \right) + e^{-\varepsilon_n H_{P_0}} (C_{P_{n,B_n}^0} - C_{P_0}) \right\}^2 \\ &= E_{B_n} P_0 \left\{ C_{P_{n,B_n}^0} \left(e^{-\varepsilon_n H_{P_0}} \varepsilon_n (H_{P_{n,B_n}^0} - H_{P_0}) + \frac{e^{Y_{P_{n,B_n}^0}}}{2} \varepsilon_n^2 (H_{P_{n,B_n}^0} - H_{P_0})^2 \right) \right. \\ &\quad \left. + e^{-\varepsilon_n H_{P_0}} (C_{P_{n,B_n}^0} - C_{P_0}) \right\}^2 \\ &= E_{B_n} P_0 C_{P_{n,B_n}^0}^2 \left(e^{-\varepsilon_n H_{P_0}} \varepsilon_n (H_{P_{n,B_n}^0} - H_{P_0}) + \frac{e^{Y_{P_{n,B_n}^0}}}{2} \varepsilon_n^2 (H_{P_{n,B_n}^0} - H_{P_0})^2 \right)^2 \\ &+ 2E_{B_n} P_0 \left\{ C_{P_{n,B_n}^0} \left(e^{-\varepsilon_n H_{P_0}} \varepsilon_n (H_{P_{n,B_n}^0} - H_{P_0}) + \frac{e^{Y_{P_{n,B_n}^0}}}{2} \varepsilon_n^2 (H_{P_{n,B_n}^0} - H_{P_0})^2 \right) \times \right. \\ &\quad \left. e^{-\varepsilon_n H_{P_0}} (C_{P_{n,B_n}^0} - C_{P_0}) \right\} \\ &+ E_{B_n} P_0 \left\{ e^{-2\varepsilon_n H_{P_0}} (C_{P_{n,B_n}^0} - C_{P_0})^2 \right\} \end{aligned}$$

$$\begin{aligned}
&= E_{B_n} P_0 C_{P_{n,B_n}^0}^2 \left(e^{-2\varepsilon_n H_{P_0}} \varepsilon_n^2 (H_{P_{n,B_n}^0} - H_{P_0})^2 \right) + E_{B_n} P_0 \frac{e^{2Y_{P_{n,B_n}^0}}}{4} \varepsilon_n^4 (H_{P_{n,B_n}^0} - H_{P_0})^4 \\
&+ 2E_{B_n} P_0 \left\{ e^{-\varepsilon_n H_{P_0}} \frac{e^{Y_{P_{n,B_n}^0}}}{2} \varepsilon_n^3 (H_{P_{n,B_n}^0} - H_{P_0})^3 \right\} \\
&+ 2E_{B_n} P_0 \left\{ C_{P_{n,B_n}^0} \left(e^{-\varepsilon_n H_{P_0}} \varepsilon_n (H_{P_{n,B_n}^0} - H_{P_0}) + \frac{e^{Y_{P_{n,B_n}^0}}}{2} \varepsilon_n^2 (H_{P_{n,B_n}^0} - H_{P_0})^2 \right) \times \right. \\
&\quad \left. e^{-\varepsilon_n H_{P_0}} (C_{P_{n,B_n}^0} - C_{P_0}) \right\} \\
&+ E_{B_n} P_0 \left\{ e^{-2\varepsilon_n H_{P_0}} (C_{P_{n,B_n}^0} - C_{P_0})^2 \right\}.
\end{aligned}$$

After repeated applications of Cauchy-Schwartz inequality to the summands, the boundedness assumptions and lemmas 3.4 and 3.9 imply that indeed

$$E_{B_n} P_0 \left(C_{P_{n,B_n}^0} e^{-\varepsilon_n H_{P_{n,B_n}^0}} - C_{P_0} e^{-\varepsilon_n H_{P_0}} \right)^2 \xrightarrow{P} 0.$$

Hence (3.43) converges to 0 in probability.

The term (3.44) can be written as

$$\begin{aligned}
&E_{B_n} (P_{n,B_n}^1 - P_0) D_Y(P_0, \varepsilon_n) \equiv E_{B_n} (P_{n,B_n}^1 - P_0) (Y - \hat{Q}(\varepsilon_n)(P_0)) H_{\hat{g}(P_0)}^* \\
&= E_{B_n} (P_{n,B_n}^1 - P_0) Y H_{\hat{g}(P_0)}^* - E_{B_n} (P_{n,B_n}^1 - P_0) \hat{Q}(\varepsilon_n)(P_0) H_{\hat{g}(P_0)}^*
\end{aligned}$$

The first summand in the last equality is an empirical difference that is asymptotically normal with mean zero. In particular, it converges to zero in probability. The second summand also converges to 0 in probability. To see that, let $\mathcal{F}(P_0) = \{f_\varepsilon = \hat{Q}(P_0)(\varepsilon) H_{\hat{g}(P_0)}^* : \varepsilon\}$, where ε ranges over K . On a sample split of B_n , lemma 3.1 implies that

$$\sqrt{n} E \left(\sup_{f \in \mathcal{F}} |(P_{n,B_n}^1 - P_0) f| \right) \leq \text{Entro}(\mathcal{F}) \sqrt{P_0 \mathbf{F}^2},$$

where \mathbf{F} is an envelope of \mathcal{F} . Therefore, we may write

$$\begin{aligned}
&P \left(|(P_{n,B_n}^1 - P_0) f_{\varepsilon_n}(P_0)| > \delta \right) \leq EP \left(\sup_{f \in \mathcal{F}} |(P_{n,B_n}^1 - P_0) f| > \delta \right) \\
&\leq \frac{1}{\delta} EE \left(\sup_{f \in \mathcal{F}} |(P_{n,B_n}^1 - P_0) f| \right) \leq \frac{1}{\sqrt{n}} \frac{1}{\delta} E \text{Entro}(\mathcal{F}) \sqrt{P_0 \mathbf{F}^2}.
\end{aligned}$$

The entropy of this class is bounded. From the boundedness assumptions of $\hat{g}(P_0)$ and the definition of $\hat{Q}(\varepsilon)$, we see that all the functions the \mathcal{F} are also bounded, hence $\sqrt{P_0 \mathbf{F}^2}$ is

bounded. Therefore, the RHS of the last inequality converges to 0 in probability as $n \rightarrow \infty$. This result combined with the fact that B_n is uniformly distributed over a finite support now imply that $E_{B_n} \left(P_{n,B_n}^1 - P_0 \right) \hat{Q}(\varepsilon_n)(P_0) H_{\hat{g}(P_0)}^*$ indeed converge to 0 in probability.

It remains to show that (3.42) converges to 0 in probability. By our assumption, there exists constant $M > 0$ such that $P(|\varepsilon_n| < M) = 1$. Conditional on P_{n,B_n}^0 , consider the class

$$\mathcal{F}(P_{n,B_n}^0) = \{f_\varepsilon(P_{n,B_n}^0) = D_Y(P_{n,B_n}^0, \varepsilon) - D_Y(P_0, \varepsilon) : |\varepsilon| < M\}.$$

Lemma 3.1 implies that

$$\sqrt{n}E \left(\sup_{f \in \mathcal{F}(P_{n,B_n}^0)} |(P_{n,B_n}^1 - P_0)f| \right) \leq Entro(\mathcal{F}(P_{n,B_n}^0)) \sqrt{P_0 \mathbf{F}(P_{n,B_n}^0)^2},$$

where $\mathbf{F}(P_{n,B_n}^0)$ is an envelope of $\mathcal{F}(P_{n,B_n}^0)$. Therefore, we may write

$$\begin{aligned} & P \left(|(P_{n,B_n}^1 - P_0)f_{\varepsilon_n}(P_{n,B_n}^0)| > \delta \right) \\ & \leq EP \left(\sup_{f \in \mathcal{F}(P_{n,B_n}^0)} |(P_{n,B_n}^1 - P_0)f| > \delta \middle| P_{n,B_n}^0 \right) \\ & \leq \frac{1}{\delta} EE \left(\sup_{f \in \mathcal{F}(P_{n,B_n}^0)} |(P_{n,B_n}^1 - P_0)f| \right) \\ & \leq \frac{1}{\sqrt{n}} \frac{1}{\delta} EE Entro(\mathcal{F}(P_{n,B_n}^0)) \sqrt{P_0 \mathbf{F}(P_{n,B_n}^0)^2}. \end{aligned}$$

Firstly note that f_ε is bounded per our assumptions and construction of $\hat{Q}(P)(\varepsilon)$. Hence $\sqrt{P_0 \mathbf{F}(P_{n,B_n}^0)^2}$ is bounded. On the other hand, the entropy of the class is also bounded. Therefore, we indeed have $P \left(|(P_{n,B_n}^1 - P_0)f_{\varepsilon_n}(P_0)| > \delta \right)$ converges to 0 as $n \rightarrow \infty$. Consequently, (3.42) converges to 0 in probability. We have thus shown that $P_0 D_Y(P_0, \varepsilon_n) \xrightarrow{P} 0$.

Since K is compact, there is a subsequence ε_{nk} such that $\varepsilon_{nk} \xrightarrow{P} \varepsilon^*$ for some $\varepsilon^* \in K$. This implies that for

$$\begin{aligned} g(\varepsilon) & \equiv P_0 D_Y(P_0, \varepsilon) \\ & = P_0 Y H_{\hat{g}(P_0)}^* - P_0 \frac{H_{\hat{g}(P_0)}^*}{1 + e^{-\text{logit}(\hat{Q}(P_0)) - \varepsilon H_{\hat{g}(P_0)}^*}}, \end{aligned}$$

which is continuous over K , we must have $g(\varepsilon_{nk}) \xrightarrow{P} g(\varepsilon^*)$.

Since $g(\varepsilon_n) \xrightarrow{P} 0$, as determined above, it follows that $g(\varepsilon^*) = 0$. On the other hand, by definition of ε_0 we have that $g(\varepsilon_0) = 0$. Note that $g'(\varepsilon) < 0$, hence it's monotonic in

ε . Therefore we indeed have $\varepsilon^* = \varepsilon_0$. This implies that all convergent subsequences of ε_n converge to ε_0 in probability. Since K is compact, it now implies that ε_n converge to ε_0 in probability. \square

Proof of Lemma 3.12:

Conditional on P_{n,B_n}^0 , for a deterministic sequence δ_n converging to 0, consider the class

$$\mathcal{F}_{\delta_n}(P_{n,B_n}^0) \equiv \left\{ D_Y(P_{n,B_n}^0, \varepsilon) - D_Y(P_0, \varepsilon) : \|\varepsilon - \varepsilon_0\| < \delta_n \right\},$$

where

$$\begin{aligned} & D_Y(P_{n,B_n}^0, \varepsilon) - D_Y(P_0, \varepsilon) \\ &= \left(Y - \hat{Q}(P_{n,B_n}^0)(\varepsilon) \right) H_{\hat{g}(P_{n,B_n}^0)}^* - \left(Y - \hat{Q}(P_0)(\varepsilon_0) \right) H_{\hat{g}(P_0)}^*. \end{aligned}$$

From lemma 3.11, we readily have $\|\varepsilon_n - \varepsilon_0\| \xrightarrow{P} 0$. To obtain the desired result, it remains to show that $\mathcal{F}_{\delta_n}(P_{n,B_n}^0)$ satisfies the conditions of lemma 3.2.

For convenience, let $H_{P_{n,B_n}^0}(O) \equiv H_{\hat{g}(P_{n,B_n}^0)}^*(A, W)$, and H_{P_0} its counterpart at P_0 . Then, we can find an envelope for this class of functions as follows:

$$\begin{aligned} & |D_Y(P_{n,B_n}^0, \varepsilon) - D_Y(P_0, \varepsilon)| \\ & \leq |Y| |H_{P_{n,B_n}^0} - H_{P_0}| + \hat{Q}(P_0)(\varepsilon_0) |H_{P_{n,B_n}^0} - H_{P_0}| + |H_{P_{n,B_n}^0}| |\hat{Q}(P_{n,B_n}^0)(\varepsilon) - \hat{Q}(P_0)(\varepsilon_0)| \\ & \leq |Y + \hat{Q}(P_0)(\varepsilon_0)| |H_{P_{n,B_n}^0} - H_{P_0}| \\ & \quad + |H_{P_{n,B_n}^0}| \left| \frac{C_{P_{n,B_n}^0} e^{-\varepsilon H_{P_{n,B_n}^0}} - C_{P_0} e^{-\varepsilon_0 H_{P_0}}}{(1 + C_{P_{n,B_n}^0} e^{-\varepsilon H_{P_{n,B_n}^0}})(1 + C_{P_{n,B_n}^0} e^{-\varepsilon H_{P_{n,B_n}^0}})} \right| \\ & \leq |Y + \hat{Q}(P_0)(\varepsilon_0)| |H_{P_{n,B_n}^0} - H_{P_0}| + |H_{P_{n,B_n}^0}| e^{-\varepsilon_0 H_{P_0}} |C_{P_{n,B_n}^0} - C_{P_0}| \\ & \quad + |H_{P_{n,B_n}^0}| |C_{P_{n,B_n}^0}| |e^{-\varepsilon H_{P_{n,B_n}^0}} - e^{-\varepsilon_0 H_{P_0}}| \\ & \leq |Y + \hat{Q}(P_0)(\varepsilon_0)| |H_{P_{n,B_n}^0} - H_{P_0}| + |H_{P_{n,B_n}^0}| e^{-\varepsilon_0 H_{P_0}} |C_{P_{n,B_n}^0} - C_{P_0}| \\ & \quad + |H_{P_{n,B_n}^0}| |C_{P_{n,B_n}^0}| |e^{-\varepsilon_0 H_{P_0}}| |\varepsilon H_{P_{n,B_n}^0} - \varepsilon_0 H_{P_0}| \\ & \quad + |H_{P_{n,B_n}^0}| |C_{P_{n,B_n}^0}| \frac{e^{M''}}{2} |\varepsilon H_{P_{n,B_n}^0} - \varepsilon_0 H_{P_0}|^2 \end{aligned}$$

$$\begin{aligned}
&\leq |Y + \hat{Q}(P_0)(\varepsilon_0)| |H_{P_{n,B_n}^0} - H_{P_0}| + |H_{P_{n,B_n}^0}| |e^{-\varepsilon_0 H_{P_0}}| |C_{P_{n,B_n}^0} - C_{P_0}| \\
&+ |H_{P_{n,B_n}^0}| |C_{P_{n,B_n}^0}| e^{-\varepsilon_0 H_{P_0}} \varepsilon_0 |H_{P_{n,B_n}^0} - H_{P_0}| \\
&+ |H_{P_{n,B_n}^0}|^2 |C_{P_{n,B_n}^0}| e^{-\varepsilon_0 H_{P_0}} \delta_n + |H_{P_{n,B_n}^0}| |C_{P_{n,B_n}^0}| \frac{e^{M''}}{2} \varepsilon_0^2 |H_{P_{n,B_n}^0} - H_{P_0}|^2 \\
&+ |H_{P_{n,B_n}^0}|^3 |C_{P_{n,B_n}^0}| \frac{e^{M''}}{2} \delta_n^2 \\
&+ 2 |H_{P_{n,B_n}^0}|^2 |C_{P_{n,B_n}^0}| \frac{e^{M''}}{2} \varepsilon_0 |H_{P_{n,B_n}^0} - H_{P_0}| \delta_n \\
&= \left(Y + \hat{Q}(P_0)(\varepsilon_0) + |H_{P_{n,B_n}^0}| |C_{P_{n,B_n}^0}| e^{-\varepsilon_0 H_{P_0}} \varepsilon_0 \right) |H_{P_{n,B_n}^0} - H_{P_0}| \\
&+ |H_{P_{n,B_n}^0}| |e^{-\varepsilon_0 H_{P_0}}| |C_{P_{n,B_n}^0} - C_{P_0}| \\
&+ |H_{P_{n,B_n}^0}|^2 |C_{P_{n,B_n}^0}| e^{-\varepsilon_0 H_{P_0}} \delta_n + |H_{P_{n,B_n}^0}| |C_{P_{n,B_n}^0}| \frac{e^{M''}}{2} \varepsilon_0^2 |H_{P_{n,B_n}^0} - H_{P_0}|^2 \\
&+ |H_{P_{n,B_n}^0}|^3 |C_{P_{n,B_n}^0}| \frac{e^{M''}}{2} \delta_n^2 \\
&+ 2 |H_{P_{n,B_n}^0}|^2 |C_{P_{n,B_n}^0}| \frac{e^{M''}}{2} \varepsilon_0 |H_{P_{n,B_n}^0} - H_{P_0}| \delta_n \\
&\equiv \mathbf{F}_n.
\end{aligned}$$

Applying Cauchy-Schwartz inequality in combination with lemma 3.10 and boundedness assumptions, we thereby have that $EP_0(\mathbf{F}_n)^2 \rightarrow 0$. Furthermore, the entropy of $\mathcal{F}_{\delta_n}(P_{n,B_n}^0)$ is bounded. Therefore, from lemma 3.2 it follows that

$$\sqrt{n}(P_{n,B_n}^1 - P_0) \{D_Y(P_{n,B_n}^0, \varepsilon) - D_Y(P_0, \varepsilon)\} = o_P(1).$$

□.

Proof of lemma 3.13: This is proved analogue to the proof of lemma 3.12. □.

Chapter 4

Bibliography

4.1 Bibliography

- A. C. Atkinson and A. Biswas. Adaptive biased-coin designs for skewing the allocation proportion in clinical trials with normal responses. *Stat. Med.*, 24(16):2477–2492, 2005.
- U. Bandyopadhyay and A. Biswas. Adaptive designs for normal responses with prognostic factors. *Biometrika*, 88(2):409–419, 2001.
- P. J. Bickel, C. A. J. Klaassen, Y. Ritov, and J. A. Wellner. *Efficient and adaptive estimation for semiparametric models*. Springer-Verlag, New York, 1998. Reprint of the 1993 original.
- A. Biswas, R. Bhattacharya, and E. Park. On a class of optimal covariate-adjusted response-adaptive designs for survival outcomes. *Statistical methods in medical research*, 2014.
- A. Chambaz and M. J. van der Laan. Targeting the optimal design in randomized clinical trials with binary outcomes and no covariate: Theoretical study. *Int. J. Biostat.*, 7(1), 2011a. Article 10.
- A. Chambaz and M. J. van der Laan. Targeting the optimal design in randomized clinical trials with binary outcomes and no covariate: Simulation study. *Int. J. Biostat.*, 7(1), 2011b. Article 11.

- A. Chambaz and M. J. van der Laan. Estimation and testing in targeted group sequential covariate-adjusted randomized clinical trials. Technical report 299, Division of Biostatistics, University of California, Berkeley, 2011c.
- A. Chambaz and M. J. van der Laan. Inference in targeted group sequential covariate-adjusted randomized clinical trials. *Scandinavian Journal of Statistics*, 41(1):104–140, 2013.
- Y. I. Chang and E. Park. Sequential estimation for covariate-adjusted response-adaptive designs. *J. Korean Statistical Society*, 42(1):105–116, 2013.
- J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22, 2010. URL <http://www.jstatsoft.org/v33/i01/>.
- S. Gruber and M.J. van der Laan. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *International Journal of Biostatistics*, 6, 2010.
- M.A. Hernan, B. Brumback, and J.M. Robins. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology*, 11(5):561–570, 2000.
- F. Hu and W. F. Rosenberger. Optimality, variability, power: evaluating response-adaptive randomization procedures for treatment comparisons. *Journal of the American Statistical Association*, 98(463):671–678, 2003.
- F. Hu and W. F. Rosenberger. *The theory of response-adaptive randomization in clinical trials*, volume 525. John Wiley & Sons, 2006.
- T. Huang, Z. Liu, and F. Hu. Longitudinal covariate-adjusted response-adaptive randomized designs. *J. Statistical Planning and Inference*, 143(10):1816–1827, 2013.
- A. Ivanova. A play-the-winner type urn model with reduced variability. *Metrika*, 58:1–13, 2003.
- C. Jennison and B. W. Turnbull. *Group Sequential Methods with Applications to Clinical Trials*. Chapman & Hall/CRC, Boca Raton, FL, 2000.
- K. Moore and M. J. van der Laan. Covariate adjustment in randomized trials with binary outcomes: targeted maximum likelihood estimation. *Stat. Med.*, 28(1):39–64, 2009.
- J. Pearl. *Causality*. Cambridge University Press, Cambridge, 2000. Models, reasoning, and inference.

- M. Petersen, K. Porter, S. Gruber, Y. Wang, and M.J. van der Laan. Diagnosing and responding to violations in the positivity assumption. Technical report 269, Division of Biostatistics, University of California, Berkeley, 2010. URL <http://www.bepress.com/ucbbiostat/paper269>.
- D. Pollard. *Convergence of stochastic processes*. Springer Series in Statistics. Springer-Verlag, New York, 1984. ISBN 0-387-90990-7.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014. URL <http://www.R-project.org/>.
- J.M. Robins. A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7:1393–1512, 1986.
- J.M. Robins. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical models in epidemiology: the environment and clinical trials*, pages 95–134. Springer-Verlag, 1999.
- J.M. Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association*, 2000.
- J.M. Robins and A. Rotnitzky. Comment on the Bickel and Kwon article, “Inference for semiparametric models: Some questions and an answer”. *Statistica Sinica*, 11(4): 920–936, 2001.
- J.M. Robins, A. Rotnitzky, and M.J. van der Laan. Comment on “on profile Likelihood” by S.A. Murphy and A.W. van der Vaart. *Journal of the American Statistical Association – Theory and Methods*, 450:431–435, 2000.
- W. F. Rosenberger. New directions in adaptive designs. *Statist. Sci.*, 11:137–149, 1996.
- W. F. Rosenberger and F. Hu. Maximizing power and minimizing treatment failures in clinical trials. *Clinical Trials*, 1(141-147), 2004.
- W. F. Rosenberger, A. N. Vidyashankar, and D. K. Agarwal. Covariate-adjusted response-adaptive designs for binary response. *Journal of biopharmaceutical statistics*, 11(4): 227–236, 2001.
- W. F. Rosenberger, O. Sverdlov, and F. Hu. Adaptive randomization for clinical trials. *J Biopharm Stat*, 22(4):719–36, 2012.

- M. Rosenblum. *Robust Analysis of RCTs Using Generalized Linear Models*, chapter 11. Springer-Verlag, New York, 2011.
- P. K. Sen and J. M. Singer. *Large sample methods in statistics*. Chapman & Hall, New York, 1993. ISBN 0-412-04221-5. An introduction with applications.
- J. Shao and X. Yu. Validity of tests under covariate-adaptive biased coin randomization and generalized linear models. *Biometrics*, 69(4):960–969, 2013.
- J. Shao, X. Yu, and B. Zhong. A theory for testing hypotheses under covariate-adaptive randomization. *Biometrika*, 97(2):347–360, 2010.
- O.M. Stitelman and M.J. van der Laan. Collaborative targeted maximum likelihood for time-to-event data. Technical Report 260, Division of Biostatistics, University of California, Berkeley, 2010.
- O. Sverdlov, W. F. Rosenberger, and Y. Ryznik. Utility of covariate-adjusted response-adaptive randomization in survival trials. *Statistics in Biopharmaceutical Research*, 5(1):38–53, 2013.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- M. J. van der Laan and J. M. Robins. *Unified methods for censored longitudinal data and causality*. Springer Series in Statistics. Springer-Verlag, New York, 2003.
- M. J. van der Laan and D. Rubin. Targeted maximum likelihood learning. *Int. J. Biostat.*, 2(1), 2006.
- M.J. van der Laan. Causal effect models for intention to treat and realistic individualized treatment rules. Technical report 203, Division of Biostatistics, University of California, Berkeley, 2006.
- M.J. van der Laan and S. Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. Technical report, Division of Biostatistics, University of California, Berkeley, November 2003.
- M.J. van der Laan and S. Gruber. Collaborative double robust penalized targeted maximum likelihood estimation. *The International Journal of Biostatistics*, 6(1), 2010.
- M.J. van der Laan, S. Rose, and S. Gruber. Readings on targeted maximum likelihood estimation. *Technical report, working paper series <http://www.bepress.com/ucbbiostat/paper254>*, September, 2009.

- A. W. van der Vaart. *Asymptotic statistics*, volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998.
- A. W. van der Vaart and J. A. Wellner. *Weak Convergence*. Springer, 1996a.
- A. W. van der Vaart and J. A. Wellner. Empirical processes indexed by estimated functions. *Lecture Notes-Monograph Series*, pages 234–252, 2007.
- A.W. van der Vaart and J.A. Wellner. *Weak Convergence and Empirical Processes*. Springer-Verlag, New York, 1996b.
- A.W. van der Vaart, S. Dudoit, and M.J. van der Laan. Oracle inequalities for multi-fold cross-validation. *Statistics and Decisions*, 24(3):351–371, 2006.
- R. van Handel. On the minimal penalty for markov order estimation. *Probability Theory and Related Fields*, 150:709–738, 2011.
- L. J. Wei and S. Durham. The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association*, 73(840-843), 1978.
- L-X. Zhang and F-F. Hu. A new family of covariate-adjusted response-adaptive designs and their properties. *Appl. Math. J. Chinese Univ. Ser. B*, 24(1):1–13, 2009.
- L-X. Zhang, F. Hu, S. H. Cheung, and W. S. Chan. Asymptotic properties of covariate-adjusted response-adaptive designs. *Ann. Statist.*, 35(3):1166–1182, 2007.