University of California, Berkeley
U.C. Berkeley Division of Biostatistics Working Paper Series

Year 2015 *Paper* 338

One-Step Targeted Minimum Loss-based Estimation Based on Universal Least Favorable One-Dimensional Submodels

Mark J. van der Laan*

*University of California, Berkeley, Division of Biostatistics, laan@berkeley.edu

One-Step Targeted Minimum Loss-based Estimation Based on Universal Least Favorable One-Dimensional Submodels

Mark J. van der Laan

Abstract

Consider a study in which one observes n independent and identically distributed random variables whose probability distribution is known to be an element of a particular statistical model, and one is concerned with estimation of a particular real valued pathwise differentiable target parameter of this data probability distribution. The canonical gradient of the pathwise derivative of the target parameter, also called the efficient influence curve, defines an asymptotically efficient estimator as an estimator that is asymptotically linear with influence curve equal to the efficient influence curve.The targeted maximum likelihood estimator is a two stage estimator obtained by constructing a so called least favorable parametric submodel through an initial estimator with score, at zero fluctuation of the initial estimator, that spans the efficient influence curve, and iteratively maximizing the corresponding parametric likelihood till no more updates occur, at which point the updated initial estimator solves the so called efficient influence curve equation. The latter property establishes the asymptotic efficiency of the TMLE under appropriate conditions, including that the initial estimator is within a neighborhood of the true data distribution.

In this article we construct a one-dimensional universal least favorable submodel for which the TMLE only takes one step, and thereby requires minimal extra fitting with data to achieve its goal of solving the efficient influence curve equation. We generalize these to universal least favorable submodels through the relevant part of the data distribution as required for targeted minimum loss-based estimation, and to universal score-specific submodels for solving any other desired equation beyond the efficient influence curve equation. We demonstrate the one-step targeted minimum loss-based estimators based on such universal least favorable submodels for a variety of examples showing that any of the goals for TMLE we previously achieved with local (typically multivariate) least favorable parametric submodels and an iterative TMLE can also be achieved with our new one-dimensional universal least favorable submodels, resulting in new one-step TMLEs for a large class of estimation problems previously addressed. Finally, remarkably, given a multidimensional target parameter, we develop a universal canonical one-dimensional submodel such that the one-step TMLE, only maximizing the log-likelihood over a univariate parameter, solves the multivariate efficient influence curve equation. This allows us to construct a one-step TMLE based on a one-dimensional parametric submodel through the initial estimator, that solves any multivariate desired set of estimating equations.

1 Introduction

Big data is taking over the world. The dimension of the data per unit as well as the number of units on which one collects data has increased drastically over time. We want to use data to learn the answer to a specific question about the underlying experiment that generated the data. Many of the current statistical methods are outdated by relying on parametric models that are much too simplistic to describe the reality behind the data, making their coefficients non-interpretable, and statistical inference more of game than a serious scientific effort. By not viewing the selection of the statistical model as a choice driven by the actual data experiment and the knowledge we have, the choice of statistical model and corresponding estimation and statistical inference becomes a non-scientific choice. By starting out with a simplistic unreasonable formulation of the actual estimation problem, one has given up on understanding the reality and one will be blinded regarding any challenges one needs to deal with in order to approximate the true answer as best as possible. An important consequence of this arbitrariness is that different data analysts that are given the same data and scientific question of interest will end up reporting drastically different results simply by selecting different models (Starmans, 2011). That is, if in the practice of statistics we do not respect the true meaning of a statistical model as a set that is known to contain the true data distribution, the field of statistics will be a wild-west: anything goes! It will lack respect for data and for the science behind it.

Targeted learning (van der Laan and Rubin, 2006; van der Laan, 2008; Rose and van der Laan, 2011) is a subfield of statistics concerned with the development of targeted machine learning algorithms that provide statistical inference for specific target parameters of the data distribution, across possible data distributions within a realistic statistical model. By necessity any such procedure cannot rely on arbitrary choices such as the selection of a parametric model or a specific machine learning algorithm, and needs to construct estimators whose sampling distribution can be estimated so that valid confidence intervals can be constructed. The latter requires the estimator to be as unbiased as possible so that the bias can be ignored in statistical inference. The first step in targeted learning is the formulation of a statistical model that contains the true probability distribution of the data, at least till close approximation, and includes the actual available knowledge about the experiment, requiring deep study of the experiment. Given such statistical models, the target parameter needs to be carefully defined to best approximate the answer to the scientific question of interest. This often involves defining an underlying causal or full-data model, defining a causal quantity or full-data target parameter of interest, and establishing an estimand that identifies this target quantity from the observed data distribution. The statistical estimation problem is now defined by the statistical model and the target parameter mapping that maps any probability distribution in the statistical model in the corresponding value of the target parameter. By having defined the estimation problem honestly, the estimand can be honestly interpreted, and any estimator can now be scientifically judged and evaluated, both theoretically and practically through simulation studies. A possible sensitivity analysis might shed some more light about the possible discrepancy between the estimand and the underlying target quantity of interest, due to violations of the identifiability assumptions.

Fortunately, efficiency (Bickel et al., 1997) and empirical process theory (van der Vaart and Wellner, 1996) for general statistical models provides a great foundation for the construction of such targeted machine learning algorithms. The canonical gradient of the pathwise derivative of the target parameter mapping defines an asymptotically efficient estimator as an estimator that is asymptotically linear with influence curve equal to the canonical gradient, which is the reason that the canonical gradient is also called the efficient influence curve. The construction of an efficient estimator of a pathwise differentiable target parameter will thereby naturally involve the utilization of this canonical gradient. The one-step estimator (e.g., (Bickel et al., 1997)) is such a general method that adds to an initial estimator of the target parameter the empirical mean of the estimated efficient influence curve. Estimating equation methodology (van der Laan and Robins, 2003; Robins and Rotnitzky, 1992) represents a related methodology that assumes that the efficient influence curve can be represented as an estimating function in the target parameter and a nuisance parameter, and defines the estimator as the solution of the resulting estimating equation. The targeted maximum likelihood estimator is a two stage estimator obtained by constructing a parametric submodel through an initial estimator of the data distribution with score, at zero fluctuation of the initial estimator, that spans the efficient influence curve, and iteratively maximizing the corresponding parametric likelihood till no more updates occur, at which point the updated initial estimator solves the so called efficient influence curve equation. The TMLE of the target parameter is now the corresponding plug-in estimator. The fact that the targeted estimator of the data distribution solves the efficient influence curve equation provides the basis for establishing the asymptotic efficiency of the TMLE under regularity conditions, beyond the crucial condition that the initial estimator is within a neighborhood (e.g., $n^{-1/4}$) of the true data distribution. To minimize the degree of violation of this crucial rate-of-convergence condition on the initial estimator as much as possible, we have proposed to construct such an initial estimator with the ensemble super-learner template fully utilizing the power and generality of cross-validation(van der Laan and Dudoit, 2003; van der Vaart et al., 2006; van der Laan et al., 2006, 2007; Polley et al., 2012), while integrating the state of the art in machine learning. This super-learner has been proven to be optimal in the sense that it performs asymptotically as well as the best weighted combination of candidate estimators in its library of candidate estimators.

This parametric submodel through the initial estimator with a score that spans the efficient influence curve is called least favorable because it is the parametric submodel that maximizes the asymptotic variance of the submodel-specific maximum likelihood estimator of the target parameter under sampling from the initial estimator. In this article, we point out that this least favorable parametric submodel can also be interpreted as the submodel that maximizes the absolute infinitesimal change in target parameter (relative to initial estimator) divided by the information-norm of the infinitesimal change in probability distribution (relative to initial estimator). This provides a nice intuition about the targeted maximum likelihood step in TMLE as a fitting procedure that locally maximizes the change in target parameter per unit amount of fitting as measured by unit of information. However, it also shows that this choice of submodel is tailored to be optimal locally around the initial estimator, so that its optimality relies on the initial estimator being close enough to the true probability distribution.

This motivates us in this article to define and construct a one-dimensional universal least favorable submodel whose score equals the efficient influence curve at *each* of its parameter values, not just at 0, and show that such a universal least favorable submodel makes the targeted maximum likelihood estimator perform the desired job in one step, with minimal additional fitting of the data. As a consequence, it maximally preserves the statistical performance of the initial estimator, while achieving its desired targeted bias reduction. In particular, this universal least favorable submodel avoids the need for iterative targeted maximum likelihood estimation, and thereby possible overfitting in finite samples. It also provides the basis to various generalizations as needed for targeted minimum loss-based estimation of a possibly multivariate target parameter.

1.1 Organization of article

Up till the last few sections we will focus on one-dimensional target parameters. In Section 2 we provide the above mentioned intuition behind a local least favorable one-dimensional parametric submodel through an initial estimator of the data distribution as a submodel that maximizes, at zero fluctuation, the rate at which the target parameter changes per change in initial estimator of data distribution, where the latter is measured by the Kullback-Leibler divergence. Motivated by this intuition, in Section 2 we define universal least favorable one-dimensional submodels, and show that these now maximize the above rate of change in target parameter at each amount of fluctuation of the initial estimator, and demonstrate that this property results in a one-step targeted maximum likelihood estimator. In Section 3 we define a general universal least favorable submodel analytically in terms of a differential and integral equation, and show that its key property indeed holds. We also present its corresponding practical implementation based on discretizing the differential equation and integral equation while making sure that, in spite of the discretization, it results in a submodel by using a local least favorable parametric submodel. In the Appendix we showcase this universal least favorable submodel and its corresponding one-step TMLE for estimation of a particular target parameter in a (high-dimensional) parametric model. In Section 4 we generalize these universal least favorable submodels for fluctuating a data distribution (defined as minimizer of the risk of the log-likelihood loss) to universal least favorable submodels fluctuating an infinite dimensional parameter of the data distribution, defined as a minimizer of the risk of a specific loss-function. In Section 5 we demonstrate this loss-function specific universal least favorable submodel and corresponding one-step TMLE for nonparametric estimation of the causal effect of treatment among the treated, resulting in a new one-step TMLE relative to the previously proposed iterative TMLE (Zheng et al., 2013). In Section C in the Appendix we generalize the universal least favorable one-dimensional submodels to universal score-specific one-dimensional submodels that can be used to update an initial estimator into an estimator solving any user supplied desired score equation, not just the efficient influence curve equation. This allows us to construct one-step TMLE based on such universal score-specific one-dimensional submodels that satisfy additional desirable properties beyond asymptotic efficiency when the initial estimator converges to the truth. In Section C We demonstrate this with various examples, revisiting previous iterative TMLEs that relied on higher dimensional local least favorable submodels. In Section D in the Appendix we generalize the universal least favorable (or desired) submodels w.r.t. some loss function to the case that the loss-function depends itself on unknown nuisance parameters, and propose corresponding one-step TMLE using these universal least favorable submodels. In Section D We demonstrate this one-step TMLE for estimation of the counterfactual mean under a dynamic multiple time-point intervention based on a longitudinal data structure, resulting in a new very simple TMLE that only requires minimizing an empirical risk over a single fluctuation parameter, while our previous TMLE would use a separate parameter for each time point (Gruber and van der Laan, 2012; Petersen et al., 2013).

In Section 6 we set as our goal to generalize these universal one-dimensional least favorable submodels targeting a one-dimensional parameter to universal least favorable submodels that target a multidimensional parameter, or more generally, whose corresponding one-step TMLE map an initial data distribution into a targeted update that solves a user supplied multidimensional (e.g., efficient influence curve) equation in the probability distribution. Remarkably, we are able to construct a one-dimensional universal canonical submodel for which the resulting one-step TMLE solves the desired multidimensional efficient influence curve equation in the data distribution. We will refer to this submodel as the canonical submodel again since it is uniquely characterized by the (now multivariate) canonical gradient. As before, we will present the analytic differential and integral equation definition of this canonical submodel as well as a practical implementation analogue in terms of a local *multidimensional* least favorable submodel. The generalization to one-dimensional universal score-specific submodels that imply that the one-step TMLE solves the desired multidimensional score equation are immediate. Our presentation also applies to infinite dimensional pathwise differentiable target parameters. In Section

7 we demonstrate that this generality of our universal canonical one-dimensional submodel yields a one-step TMLE of a dynamic treatment specific counterfactual survival function, solving an outstanding problem in the literature: i.e. how can one compute an efficient substitution estimator (i.e., TMLE) of a high or even infinite dimensional target parameter. In Section 8 we present the generalization of this universal canonical submodel to arbitrary loss functions, not just the log-likelihood loss, and in Section 9 we further generalize it to loss functions that depend on a nuisance parameter. In Section 9 we also apply the corresponding one-step TMLE to obtain a new TMLE for the d-dimensional parameter of a working marginal structural model for the conditional counterfactual mean given effect modifiers of interest under a class of multiple time-point dynamic interventions. Our previous TMLE would use a least favorable submodel with a d-dimensional parameter at each time point, while this TMLE only maximizes over a single univariate parameter. Finally, we conclude with a conclusion in Section 10.

1.2 Statistical formulation of the goal and result of this article

Let $O_1, \ldots, O_n$ be n independent and identically distributed copies of a random variable $O \sim P_0$ with probability distribution P_0 that is known to be an element of a set $\mathcal{M}$ of possible probability distributions. We refer to $\mathcal{M}$ as the statistical model for the true data distribution P_0. Let $\Psi : \mathcal{M} \to \mathbb{R}^d$ be a d-dimensional target parameter mapping, so that $\psi_0 = \Psi(P_0)$ represents the target parameter or estimand of interest that best approximates the answer to the question of interest. We assume that Ψ is pathwise differentiable at each $P \in \mathcal{M}$ with canonical gradient $D^*(P)$. We will use notation $Pf = \int f(o)dP(o)$ for the expectation operator w.r.t. P. That is, for each path $\{P_{\epsilon,h} : \epsilon\}$ through P at $\epsilon = 0$ and score S_h, indexed by h in some index set $\mathcal{H}$, we have

$$\left.\frac{d}{d\epsilon}\Psi(P_{\epsilon,h})\right|_{\epsilon=0} = PD^*(P)S_h,$$

and $D^*(P)$ is the unique gradient that is also an element of the so called tangent space $T(P)$, defined as the closure of the linear span of all scores $\{S_h : h \in \mathcal{H}\}$ in the Hlibert space $L^2_0(P)$ of functions of O with mean zero under P, endowed with the inner-product $\langle S_1, S_2 \rangle = PS_1S_2$.

An estimator of ψ_0 is a mapping $\hat{\Psi}$ that maps the empirical probability distribution P_n of $O_1, \ldots, O_n$ into the parameter space $\Psi(\mathcal{M}) \subset \mathbb{R}^d$, and the corresponding estimate of ψ_0 is given by $\psi_n = \hat{\Psi}(P_n)$. An estimator $\hat{\Psi}(P_n)$ is asymptotically efficient at P_0 if and only if it is asymptotically linear with influence curve equal to the canonical gradient $D^*(P_0)$:

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

Such an estimator satisfies (by CLT) satisfies $\sqrt{n}(\psi_n - \psi_0) \Rightarrow_d N(0, \Sigma_0 = P_0\{D^*(P_0)D^*(P_0)^\top\})$, so that statistical inference can be based on the estimator of its influence curve

$D * (P_0)$. The canonical gradient $D^*(P_0)$ of $\Psi : \mathcal{M} \rightarrow \mathbb{R}^d$ is also called the efficient influence curve.

A targeted maximum likelihood estimator (TMLE) is defined as follows. One first constructs an initial estimator $P_n^0 \in \mathcal{M}$ of P_0. In addition, one defines a local least favorable parametric submodel $\{P_{n,\delta}^{0,lfm} : \delta\}$ through P_n^0 at $\delta = 0$ with d-dimensional parameter δ and with score $\frac{d}{d\delta} \log dP_{n,\delta}^{0,lfm}/dP_n^0\Big|_{\delta=0} = D^*(P_n^0)$. This is used to define the corresponding maximum likelihood estimator $\delta^0 = \arg\max P_n \log dP_{n,\delta}^{0,lfm}/dP_n^0$. The one-step TMLE of P_0 is now defined as $P_n^1 = P_{n,\delta^0}^{0,lfm}$. This process is iterated by defining $P_n^{k+1} = P_{n,\delta^k}^{k,lfm}$, $k = 1, 2, \ldots$, till a $k = K$ for which $\delta^K \approx 0$. The TMLE of P_0 is then defined by the final update $P_n^* = P_{n,\delta^K}^{K,lfm}$, which solves $P_n D^*(P_n^*) \approx 0$. The TMLE of ψ_0 is the corresponding plug-in estimator $\Psi(P_n^*)$. Here ≈ 0 can be replaced by $o_P(1/\sqrt{n})$: for example, one might iterate till $\parallel P_n D^*(P_n^K) \parallel \leq 1/n$, where one could use the Euclidean norm.

The asymptotic efficiency of the TMLE, under regularity conditions, is established as follows. First, define $R_2(P, P_0)$ by the equation $\Psi(P) - \Psi(P_0) = (P - P_0)D^*(P) + R_2(P, P_0)$, where, due to $D^*(P)$ being a canonical gradient, $R_2(P, P_0)$ will be a second order difference between P and P_0. Applying this identity to $P = P_n^*$, and using that $P_n D^*(P_n^*) = 0$, results in the identity:

$$\Psi(P_n^*) - \Psi(P_0) = (P_n - P_0)D^*(P_n^*) + R_2(P_n^*, P_0).$$

Assuming $R_2(P_n^*, P_0) = o_P(1/\sqrt{n})$, $D^*(P_n^*)$ falls with probability tending to one in a P_0-Donsker class, and $P_0\{D^*(P_n^*) - D^*(P_0)\}^2 \rightarrow 0$ in probability as $n \rightarrow \infty$, implies now the asymptotic efficiency of the substitution estimator $\Psi(P_n^*)$.

In addition, TMLE has been generalized to targeted minimum loss-based estimation (still denoted with TMLE) in which P is replaced by a $Q(P) = \arg\min_Q PL(Q)$ for some loss-function $L(Q)(O)$, $\Psi(P) = \Psi_1(Q(P))$ for some Ψ_1, $D^*(P) = D^*(Q(P), G(P))$ for some nuisance parameter G, resulting in a TMLE (Q_n^*, G_n^*) solving $P_n D^*(Q_n^*, G_n^*) = 0$, and TMLE $\Psi_1(Q_n^*)$ of ψ_0, where often $G_n^* = G_n^0$ is not updated.

In general, TMLE presents an iterative algorithm, utilizing a local parametric submodel with loss-function specific score equal to a user supplied $D()$, that maps an initial estimator $P_n^0 \in \mathcal{M}$, or an initial estimator (Q_n^0, G_n^0) of (Q_0, G_0), into an updated P_n^*, or (Q_n^*, G_n^*), with improved empirical fit w.r.t. the loss-function of P_0 or (Q_0, G_0), so that $P_n D(P_n^*) = 0$, or $P_n D(Q_n^*, G_n^*) = 0$. Due to this generality, its statistical applications are diverse and widespread, going beyond the construction of an efficient estimator of a pathwise differentiable target parameter for arbitrary semi-parametric models and pathwise differentiable target parameter mappings: collaborative targeted maximum likelihood estimation (CTMLE) for targeted estimation of the nuisance parameter in the canonical gradient (Rose and van der Laan, 2011; van der Laan and Gruber, 2010; Gruber and van der Laan, 2012; Stitelman and van der Laan, 2010; Gruber and van der Laan, 2010); cross-validated TMLE

(CV-TMLE) to robustify the bias-reduction of the TMLE-step (Zheng and van der Laan, 2011; Rose and van der Laan, 2011); guaranteed improvement w.r.t. a user supplied asymptotically linear estimator (Gruber and van der Laan, 2012; Lendle et al., 2013); targeted initial estimator through empirical efficiency maximization (Rubin and van der Laan, 2008; Rose and van der Laan, 2011); double robust inference by targeting censoring/treatment mechanism (van der Laan, 2012a); CV-TMLE to estimate data adaptive target parameters such as the risk of a candidate estimator and thereby develop a super-learner that uses CV-TMLE instead of the normal cross-validated empirical risk (van der Laan and Petersen, 2012; Díaz and van der Laan, 2013, In press); higher-order TMLE in order to replace in the above proof $R_2()$ by a higher order term (Carone et al., 2014; Diaz et al., 2015).

In particular, in order to preserve asymptotic linearity of $\Psi(P_n^*)$ with a known and desired influence curve when P_n^* is misspecified in the sense that $D^*(P_n^*)$ converges to a $D^*(P^*)$ with a $P^* \neq P_0$ that still satisfies $R_2(P^*, P_0) = 0$, it has been shown that replacing $D^*(P)$ by an appropriate score $\tilde{D}(P)$ in the tangent space at P that still satisfies $\tilde{D}(P_0) = D^*(P_0)$ is required (van der Laan, 2012b; Gruber and van der Laan, 2012). In that case, the above formulation requires a local parametric submodel of $\mathcal{M}$ through P_n^0 that generates this $\tilde{D}(P_n^0)$ as score at $\delta = 0$.

Even though the TMLE framework has been shown to be flexible enough to handle any of the challenges we have encountered, in many cases the proposed TMLE is iterative and uses a local parametric submodel through the initial estimator that has more, and possibly many more, than d (fluctuation) parameters. This can result in a small sample issues regarding convergence of the TMLE algorithm or causes finite sample instability of the estimator. It also contrast the principle goal of TMLE as being a procedure that updates the initial estimator with *minimal* extra fitting into a new estimator that solves the desired estimating equation that provides the basis of the desired asymptotic linearity and normality of the TMLE. By using an over-parameterized local submodel or by using an iterative algorithm these TMLE use more fitting of the data than should be needed to achieve the desired goal.

Goal of article: The goal set out in this article is to construct a parametric submodel $\{P_{n,\epsilon}^0 : \epsilon\}$ through an initial $P_n^0 \in \mathcal{M}$ so that the above TMLE algorithm only takes one step, and the dimension of ϵ is smaller or equal than d. The construction of this parametric submodel will be philosophically grounded by being in a sense the shortest path (with distance measured by information/data fitting needed) towards its goal (solving the desired score equation). We will first consider the case $d = 1$ and construct a one-dimensional parametric submodel satisfying this key property so that the TMLE is a one-step TMLE. We will generalize it to targeted minimum loss-based estimation, with all its variations in choice of loss function, and demonstrate it with various examples. Finally, we consider the general case $d > 1$, and construct a *one*-dimensional parametric submodel through P_n^0 for which the one-step TMLE solves each of the d desired score equations. Apparently, this one-dimensional path provides a "shortest" path towards its d-dimensional goal.

2 Intuition of TMLE: local and universal least favorable submodels

Let's consider one-dimensional target parameters (i.e., $d = 1$). A least favorable model at P is a model $\mathcal{S}^* = \{P_{\epsilon,h^*} : \epsilon\}$, dominated by P, for which $P_{\epsilon=0,h^*} = P$, and that maximizes the submodel specific Cramer-Rao lower bound defined by

$$CR(h \mid P) \equiv \frac{\left(\frac{d}{d\epsilon} \Psi(P_{\epsilon,h}) \big|_{\epsilon=0} \right)^2}{-P \frac{d^2}{d\epsilon^2} \log \frac{dP_{\epsilon,h}}{dP} \Big|_{\epsilon=0}}$$

over all such parametric submodels $\{P_{\epsilon,h} : \epsilon\}$ with h varying over some index set whose closure of the linear span generated the full tangent space $T(P) \subset L^2_0(P)$ of the model at P. Given the pathwise differentiability with canonical gradient $D^*(P)$, denoting the score of $\{P_{\epsilon,h} : \epsilon\}$ at $\epsilon = 0$ with S_h, it follows that this criterion for a submodel can be represented as follows:

$$CR(h \mid P) = \frac{(PD^*(P)S_h)^2}{PS_h^2},$$

By Cauchy-Schwarz inequality, it follows that this is maximized over all scores in the tangent space $T(P)$ by $S = D^*(P)$. Thus, a least favorable model can also be defined as any parametric model through P that has a score at P equal to $D^*(P)$.

Under some smoothness assumptions on the submodels, the criterion can also be represented as

$$CR(h \mid P) = \lim_{\epsilon \to 0} \frac{(\Psi(P_\epsilon) - \Psi(P))^2}{-2P \log dP_\epsilon / dP},$$

showing that it equals the square change in target parameter divided by the change in log-likelihood at P at an infinitesimal ϵ. Therefore, we will say that the path $\{P_{\epsilon,h^*} : \epsilon\}$ that maximizes $CR(h \mid P)$ follows at $\epsilon = 0$ (i.e., locally) a path of maximal change in target parameter per unit of information. To stress that the desired optimality property only applies locally, we will refer to such a submodel as a *locally* (i.e., at $\epsilon = 0$) least favorable submodel.

This latter representation of the criterion is intuitively appealing, since a sensible goal of a submodel $\{P_\epsilon : \epsilon\}$ through P is that a small fluctuation of P yields a maximal change in target parameter, making the MLE $\epsilon_n = \arg\max_\epsilon P_n \log dP_\epsilon/dP$ (as used in TMLE) for this parametric model locally all about fitting the target parameter, not wasting data for anything else.

The intuition of TMLE has always been to minimally increase the empirical fit of the initial estimator while achieving the desired bias reduction for the target parameter, measured by solving $P_n D^*(P_n^*)$ with a good estimator P_n^* of P_0 (so not worse than P_n^0). However, if P_n^0 is far away from P_0, the MLE ϵ_n^0 will be far from local. Even though it moves in the right direction at $\epsilon \approx 0$, there is no guarantee that it follows a path of maximal change in target parameter per change in distribution once ϵ moves farther away from zero. In the end that means that the targeted maximum likelihood estimator might not have followed such a targeted path after all, and it might have taken various iterations to finally end up with a local $\epsilon_n^K \approx 0$ at which point the algorithm stops. The distribution P_n^0 might have changed much more than needed to obtain the bias reduction in the target parameter. That is, the desired bias reduction came at an unnecessary cost of data fitting so that $\Psi(P_n^*)$ will have larger finite sample variance than needed. Based on this insight, we like to construct TMLEs that is based on a path that at each ϵ (not just at $\epsilon = 0$) follows a path of maximal change in target parameter per unit of information. We will refer to such a path as a *universal* least favorable submodel.

Definition 1 *Suppose that, given a $P \in \mathcal{M}$, $ULFM(P) = \{P_\epsilon : \epsilon \in (-a,a)\} \subset \mathcal{M}$ is a parametric submodel dominated by P, such that $P_{\epsilon=0} = P$ and for each $\epsilon \in (-a,a) \subset \mathbb{R}$, we have*

$$\frac{d}{d\epsilon} \log \frac{dP_\epsilon}{dP} = D^*(P_\epsilon). \tag{1}$$

Then, we say that $ULFM(P)$ is a universal least favorable submodel through P.

That is, this least favorable model is not only least favorable at $\epsilon = 0$, it is a least favorable model at each $P_\epsilon \in ULFM(P)$. This article proposes such universal least favorable submodels and corresponding targeted maximum likelihood and targeted minimum loss-based estimators. A very nice by-product of these universal least favorable submodels is that the TMLE always "converges" in one-step. This reflects the above intuition of a universal least favorable submodel as a shortest path submodel in the sense that it achieves the desired bias reduction at minimal increase in empirical log-likelihood.

3 A universal least favorable submodel for targeted maximum likelihood estimation

3.1 The TMLE based on a universal least favorable submodel takes only one step

Let $O_1, \ldots, O_n \sim_{iid} P_0 \in \mathcal{M}$ and $\Psi : \mathcal{M} \rightarrow \mathbb{R}$ is a one-dimensional pathwise differentiable target parameter mapping. Let $D^*(P)$ be the canonical gradient of the pathwise derivative at $P \in \mathcal{M}$ Let P_n^0 be an initial estimator of P_0. Suppose that, given a $P \in \mathcal{M}$, we can construct a universal least favorable parametric model $ULFM(P) = \{P_\epsilon : \epsilon \in (-a,a)\} \subset \mathcal{M}$. If we use this as parametric submodel in the TMLE, then the TMLE converges in one step. That is, let

$$\epsilon_n^0 = \arg\max_\epsilon P_n \log \frac{dP_{n,\epsilon}^0}{dP_n^0}.$$

One can replace the maximum ϵ_n^0 by the local maximum closest to $\epsilon = 0$, which is what we recommend in case the selected universal least favorable submodel allows for multiple local maxima. Let $P_n^1 = P^0_{n,\epsilon_n^0}$. Since ϵ_n^0 is a local maximum it solves its score equation, given by $P_n D^*(P_n^1) = 0$. That is, it achieved the goal of solving the desired efficient influence curve equation in one-step. Further iteration will not yield further updates: the next MLE

$$\epsilon_n^1 = \arg\max_\epsilon P_n \log \frac{dP^1_{n,\epsilon}}{dP_n^1} = 0.$$

Therefore, the TMLE of $\psi_0 = \Psi(P_0)$ is given by the one-step TMLE $\psi_n^* = \Psi(P_n^1)$.

In addition, we strongly suspect that a TMLE using such a least favorable model will often perform better in finite samples, certainly in situations in which the TMLE requires an iterative algorithm. In addition, it is philosophically superior by always following a path along ϵ in which the rate of square change in the parameter by unit of information is maximized at each ϵ-value.

3.2 An analytic formula for a universal least favorable submodel

This motivates us to consider if such a universal least favorable model exists and can be constructed. The answer is, yes, as our constructions below demonstrate.

In the following we use p_ϵ for the density of P_ϵ w.r.t. P, so that $p = 1$, but we will still use p (in case, one wants to use the formulas for densities w.r.t. another dominating measure). For $\epsilon \geq 0$, we recursively define

$$p_\epsilon = p \exp\left(\int_0^\epsilon D^*(P_x)dx\right), \quad (2)$$

and, for $\epsilon < 0$, we recursively define

$$p_\epsilon = p \exp\left(-\int_\epsilon^0 D^*(P_x)dx\right).$$

Theorem 1 *Consider the definition of* $\{P_\epsilon : \epsilon \in (-a, a)\}$ *above. We have that* $\{P_\epsilon : \epsilon \in (-a, a)\}$ *is a set of probability distributions dominated by* P, $P_{\epsilon=0} = P$, *and, for each* $\epsilon \in (-a, a)$, *we have*

$$\frac{d}{d\epsilon} \log \frac{dP_\epsilon}{dP} = D^*(P_\epsilon).$$

Proof: To start with it follows trivially that for for each ϵ $\frac{d}{d\epsilon} \log p_\epsilon = D^*(P_\epsilon)$. So it remains to verify that p_ϵ satisfies $\int p_\epsilon(o)dP(0) = 1$ (obviously, $p_\epsilon \geq 0$). Define $C(\epsilon, P) \equiv \int p_\epsilon dP$. Consider the probability density $p_{\epsilon,1} = C(\epsilon, P)^{-1} p_\epsilon$. We have that its score at ϵ is given by:

$$S(\epsilon, P) = \frac{1}{C(\epsilon, P)} \frac{d}{d\epsilon} C(\epsilon, P) + D^*(P_\epsilon).$$

We know that $P_\epsilon S(\epsilon, P) = 0$. Since $P_\epsilon D^*(P_\epsilon) = 0$, this implies that $\frac{d}{d\epsilon} C(\epsilon, P) = 0$. Thus, $C(\epsilon, P) = C(0, P) = 1$. This completes the proof. □

Note that this recursive relation (2) allows one to recursively solve for $p_{\epsilon+d\epsilon}$, given $\{p_x : x \in [0, \epsilon]\}$, in the sense that (e.g.) for $\epsilon > 0$,

$$\frac{p_{\epsilon+d\epsilon}}{p_\epsilon} = \exp(D^*(P_\epsilon)d\epsilon) = (1 + d\epsilon D^*(P_\epsilon)).$$

This implies a practical construction that starts with $p_{x_0=0} = p$ and recursively solves for

$$p_{x_j} = p_{x_{j-1}}(1 + (x_j - x_{j-1})D^*(P_{x_{j-1}})),\ j = 1, \ldots, N$$

for an arbitrary fine grid $0 = x_0 < x_1 < \ldots < x_N = a$. Similarly, one determines recursively

$$p_{-x_j} = p_{-x_{j-1}}(1 - (x_j - x_{j-1})D^*(P_{-x_{j-1}})),\ j = 1, \ldots, N.$$

If the model $\mathcal{M}$ is nonparametric, then this practical construction is a submodel of $\mathcal{M}$, but if the model is restricted the practical construction above might select probability distributions P_{x_j} that are not an element of $\mathcal{M}$, even though it has score at x_j equal to $D^*(P_{x_j})$ in the tangent space at P_{x_j} of the model $\mathcal{M}$. Nonetheless, this practical construction of this least favorable model can be used for any model $\mathcal{M}$ as long as one can extend the target parameter Ψ to be well defined on the probability distributions in this discrete approximation of the theoretical least favorable model, and the TMLE will still only require one step and be asymptotically efficient for the actual model $\mathcal{M}$ under regularity conditions. In addition, in the next subsection Theorem 2 proves that under mild regularity conditions, quite surprisingly, the theoretical formula (2) for this universal least favorable model, defined as a limit of the above practical construction when the partitioning gets finer and finer, is an actual submodel of $\mathcal{M}$. Another way of viewing this result is that by selecting the partitioning fine enough in the above practical construction $\{p_{x_j}, p_{-x_j} : j = 0, \ldots, N\}$ this submodel will be arbitrarily close to the model $\mathcal{M}$. Below we will also provide an alternative to the above practical construction that does preserve the submodel property while it still approximates the theoretical formula (2).

3.3 A universal least favorable submodel in terms of a local least favorable submodel

An alternative representation of the above analytic formula (2) is given by a product integral representation. Let $d\epsilon > 0$. For $\epsilon \geq 0$, we define

$$p_{\epsilon+d\epsilon} = p \prod_{x \in (0,\epsilon]} (1 + D^*(P_x)dx),$$

and for $\epsilon < 0$, we define

$$p_{\epsilon-d\epsilon} = p \prod_{x \in [\epsilon,0)} (1 - D^*(P_x)dx).$$

In other words, $p_{x+dx} = p_x(1 + D^*(P_x)dx)$, or, another way of thinking about this is that p_{x+dx} is obtained by constructing a least favorable model through P_x with score $D^*(P_x)$ at P_x, and evaluate it at parameter value dx, slightly away from zero. This suggests the following generalization of the universal least favorable model whose practical analogue will now still be an actual submodel of $\mathcal{M}$.

Let $0 = x_0 < x_1 < \ldots \leq x_N = \tau$ be an equally spaced fine grid for the interval $[0, \tau]$. Let $h = x_j - x_{j-1}$ be the width of the partition elements. We will provide a construction for P_{x_j}, $j = 0, \ldots, N$. This construction is expressed in terms of a mapping $P \rightarrow \{P_\delta^{lfm} : \delta \in (-a, a)\} \subset \mathcal{M}$ that maps any $P \in \mathcal{M}$ into a local least favorable submodel of $\mathcal{M}$ through P at $\delta = 0$ and with score $D^*(P)$ at $\delta = 0$, where a is some positive number. For any estimation problem defined by $\mathcal{M}$ and Ψ one is typically able to construct such a local least favorable submodel, so that this is hardly an assumption. Let $P_{x=0} = P$. Let $p_{x_1} = p_{x_0,h}^{lfm}$, and, in general, let $p_{x_{j+1}} = p_{x_j,h}^{lfm}$, $j = 1, 2, \ldots, N-1$. Similarly, let $-\tau = -x_N < -x_{N-1} < \ldots < -x_1 < x_0 = 0$ be the corresponding grid for $[-\tau, 0]$, and we define $p_{-x_{j+1}} = p_{-x_j,-h}^{lfm}$, $j = 1, \ldots, N-1$. In this manner, we have defined P_{x_j}, P_{-x_j}, $j = 0, \ldots, N$, and, by construction, each of these are probability distributions in the model $\mathcal{M}$. This construction is all we need when using the universal least favorable submodel in practice, such as in the TMLE.

This practical construction implies a theoretical formulation by letting N converge to infinity (i.e., let the width of the partitioning converge to zero). That is, an analytic way of representing this universal least favorable submodel, given the local least favorable model parameterization $(\epsilon, P) \rightarrow P_\epsilon^{lfm}$, is given by: for $\epsilon > 0$ and $d\epsilon > 0$, we have

$$p_{\epsilon+d\epsilon} = p_{\epsilon,d\epsilon}^{lfm}.$$

This allows for the recursive solving for p_ϵ starting at $p_{\epsilon=0} = p$, and since $p_{\epsilon,h}^{lfm} \in \mathcal{M}$, its practical approximation will never leave the model $\mathcal{M}$.

Utilizing that the least favorable model $h \rightarrow p_{\epsilon,h}^{lfm}$ is continuously twice differentiable with a score $D^*(P_\epsilon)$ at $h = 0$, we obtain a second order Tailor expansion

$$p_{\epsilon,d\epsilon}^{lfm} = p_\epsilon + \frac{d}{dh} p_{\epsilon,h}^{lfm} \bigg|_{h=0} d\epsilon + O((d\epsilon)^2) = p_\epsilon(1 + d\epsilon D^*(P_\epsilon)) + O((d\epsilon)^2),$$

so that we obtain

$$p_{\epsilon+d\epsilon} = p_\epsilon(1 + d\epsilon D^*(P_\epsilon)) + O((d\epsilon)^2).$$

This implies:

$$p_\epsilon = p \exp\left(\int_0^\epsilon D^*(P_x)dx\right).$$

Thus, we obtained the exact same representation (2) as above. This proves that, under mild regularity conditions, this analytic representation (2) is a submodel of $\mathcal{M}$ after all, but, when using its practical implementation and approximation, one should use an actual local least favorable submodel in order to guarantee that one stays in the model. We formalize this result in the following theorem.

Theorem 2 *Let $\mathcal{O}$ be a maximal support so that the support of a $P \in \mathcal{M}$ is a subset of $\mathcal{O}$. Suppose there exists a mapping $P \rightarrow \{P_\delta^{lfm} : \delta \in (-a, a)\} \subset \mathcal{M}$ that maps any $P \in \mathcal{M}$ into a local least favorable submodel of $\mathcal{M}$ through P at $\delta = 0$ and with score $D^*(P)$ at $\delta = 0$, where a is some positive number independent of P. In addition, assume the following type of second order Tailor expansion:*

$$p_{\epsilon,d\epsilon}^{lfm} = p_\epsilon + \frac{d}{dh} p_{\epsilon,h}^{lfm}\Big|_{h=0} d\epsilon + R_2(p_\epsilon, d\epsilon),$$

where

$$\sup_\epsilon \sup_{o\in\mathcal{O}} \mid R_2(p_\epsilon, d\epsilon)(o) \mid = O((d\epsilon)^2).$$

We also assume that $\sup_\epsilon \sup_{o\in\mathcal{O}} \mid D^*(P_\epsilon)p_\epsilon \mid (o) < \infty$.

Then, the universal least favorable $\{p_\epsilon : \epsilon\}$ defined by (2) is an actual submodel of $\mathcal{M}$. Its definition corresponds with $p_{\epsilon+d\epsilon} = p_{\epsilon,d\epsilon}^{lfm}$ whose corresponding practical approximation will still be a submodel.

We refer to the Appendix for an application of the universal least favorable submodel and a corresponding one-step TMLE for high dimensional parametric models.

4 Universal least favorable model for targeted minimum loss-based estimation

4.1 A universal least favorable submodel w.r.t. specific loss-function

Let's now generalize this construction of a universal least favorable w.r..t log-likelihood loss to general loss-functions so that the resulting universal least favorable submodels can be used in the more general targeted minimum loss based estimation methodology. We now assume that $\Psi(P) = \Psi_1(Q(P))$ for some parameter $Q : \mathcal{M} \rightarrow Q(\mathcal{M})$ defined on the model and real valued function Ψ_1. Let $Q(\mathcal{M}) = \{Q(P) : P \in \mathcal{M}\}$ be the parameter space of this parameter. Let $L(Q)(O)$ be a loss-function for $Q(P)$ in the sense that $Q(P) = \arg\min_{Q\in Q(\mathcal{M})} PL(Q)$. Let $D^*(P) = D^*(Q(P), G(P))$ be the canonical gradient of Ψ at P, where $G : \mathcal{M} \rightarrow G(\mathcal{M})$ is some nuisance parameter. We consider the case that the efficient influence curve is in the tangent space of Q, so that a least favorable submodel does not need to fluctuate G: otherwise, just include G in the definition of Q. Given, (Q, G), let $\{Q_\epsilon^{lfm} : \epsilon \in (-a, a)\} \subset Q(\mathcal{M})$ be a local least favorable model w.r.t. loss function $L(Q)$ at $\epsilon = 0$ so that

$$\frac{d}{d\epsilon} L(Q_\epsilon^{lfm})\Big|_{\epsilon=0} = D^*(Q, G).$$

The dependence of this submodel on G is suppressed in this notation.

Let $0 = x_0 < x_1 < \ldots < x_N = \tau$ be an equally spaced fine grid for the interval $[0, \tau]$. Let $h = x_j - x_{j-1}$ be the width of the partition elements. We present a construction for Q_{x_j}, $j = 0, \ldots, N$. Let $Q_{x=0} = Q$. Let $Q_{x_1} = Q^{lfm}_{x_0,h}$, and, in general, let $Q_{x_{j+1}} = Q^{lfm}_{x_j,h}$, $j = 1, 2, \ldots, N-1$. Similarly, let $-\tau = -x_N < -x_{N-1} < \ldots < -x_1 < x_0 = 0$ be the corresponding grid for $[-\tau, 0]$, and we define $Q_{-x_{j+1}} = Q^{lfm}_{-x_j,-h}$, $j = 1, \ldots, N-1$. In this manner, we have defined Q_{x_j}, Q_{-x_j}, $j = 0, \ldots, N$, and, by construction, each of these are an element of the parameter space $Q(\mathcal{M})$. This construction is all we need when using this submodel in practice, such as in the TMLE.

An analytic way of representing this loss-function specific universal least favorable submodel for $\epsilon \geq 0$ (and similarly for $\epsilon < 0$) is given by: for $\epsilon > 0$, $d\epsilon > 0$,

$$Q_{\epsilon+d\epsilon} = Q^{lfm}_{\epsilon,d\epsilon}, \tag{3}$$

allowing for the recursive solving for Q_ϵ starting at $Q_{\epsilon=0} = Q$, and since $Q^{lfm}_{\epsilon,h} \in Q(\mathcal{M})$, its practical approximation never leaves the parameter space $Q(\mathcal{M})$ for Q.

Let's now derive a corresponding integral equation. Assume that for some $\dot{L}(Q)(O)$, we have

$$\frac{d}{dh} L(Q^{lfm}_{\epsilon,h}) \Big|_{h=0} = \dot{L}(Q_\epsilon) \, \frac{d}{dh} Q^{lfm}_{\epsilon,h} \Big|_{h=0}.$$

Then,

$$\frac{d}{dh} Q^{lfm}_{\epsilon,h} \Big|_{h=0} = \frac{D^*(Q_\epsilon, G)}{\dot{L}(Q_\epsilon)}.$$

Utilizing that the local least favorable model $h \rightarrow Q^{lfm}_{\epsilon,h}$ is twice continuously differentiable with derivative $D^*(Q_\epsilon, G)/\dot{L}(Q_\epsilon)$ at $h = 0$, we obtain the following second order Tailor expansion:

$$\begin{aligned} Q^{lfm}_{\epsilon,d\epsilon} &= Q_\epsilon + \frac{d}{dh} Q^{lfm}_{\epsilon,h} \Big|_{h=0} d\epsilon + O((d\epsilon)^2) \\ &= Q_\epsilon + \frac{D^*(Q_\epsilon, G)}{\dot{L}(Q_\epsilon)} d\epsilon + O((d\epsilon)^2). \end{aligned}$$

This implies the following recursive analytic definition of the universal least favorable model through Q:

$$Q_\epsilon = Q + \int_0^\epsilon \frac{D^*(Q_x, G)}{\dot{L}(Q_x)} dx. \tag{4}$$

Similarly, for $\epsilon < 0$, we obtain

$$Q_\epsilon = Q - \int_\epsilon^0 \frac{D^*(Q_x, G)}{\dot{L}(Q_x)} dx.$$

As with the log-likelihood loss (and thus $Q(P)=P$), this shows that, under regularity conditions, this analytic representation for Q_ϵ is an element in $Q(\mathcal{M})$, although using it in a practical construction (in which integrals are replaced by sums) might easily leave the model space $Q(\mathcal{M})$, so that our above practical construction in terms of the local least favorable model and discrete grid represents the desired practical implementation of this universal least favorable submodel. The following theorem formalizes this result stating that the analytic formulation (4) is indeed a universal least favorable submodel.

Theorem 3 *Given, any (Q,G) compatible with model $\mathcal{M}$, let $\{Q_\delta^{lfm}:\delta\in(-a,a)\}\subset Q(\mathcal{M})$ be a local least favorable model w.r.t. loss function $L(Q)$ at $\delta=0$ so that*

$$\left.\frac{d}{d\delta}L(Q_\delta^{lfm})\right|_{\delta=0}=D^*(Q,G).$$

Assume that for some $\dot{L}(Q)(O)$, we have

$$\left.\frac{d}{d\epsilon}L(Q_\epsilon^{lfm})\right|_{\epsilon=0}=\dot{L}(Q)\left.\frac{d}{d\epsilon}Q_\epsilon^{lfm}\right|_{\epsilon=0}.$$

Consider the corresponding model $\{Q_\epsilon:\epsilon\}$ defined by (4). It goes through Q at $\epsilon=0$, and, it satisfies that for all ϵ

$$\frac{d}{d\epsilon}L(Q_\epsilon)=D^*(Q_\epsilon,G). \tag{5}$$

In addition, suppose that the $a>0$ in the local least-favorable submodel above can be chosen to be independent of the choice $(Q,G)\in\{Q_\epsilon,G_\epsilon:\epsilon\}$, and assume the following second order Tailor expansion:

$$\begin{aligned}Q_{\epsilon,d\epsilon}^{lfm}&=Q_\epsilon+\left.\frac{d}{dh}Q_{\epsilon,h}^{lfm}\right|_{h=0}d\epsilon+R_2(Q_\epsilon,G,d\epsilon)\\&=Q_\epsilon+\frac{D^*(Q_\epsilon,G)}{\dot{L}(Q_\epsilon)}d\epsilon+R_2(Q_\epsilon,G,d\epsilon),\end{aligned}$$

where

$$\sup_\epsilon\sup_{o\in\mathcal{O}}\mid R_2(Q_\epsilon,G,d\epsilon)(o)\mid=O((d\epsilon)^2).$$

We also assume that $\sup_\epsilon\sup_{o\in\mathcal{O}}\frac{|D^*(Q_\epsilon,G)}{\dot{L}(Q_\epsilon)}(o)\mid<\infty$.

Then, we also have $\{Q_\epsilon:\epsilon\}\subset Q(\mathcal{M})$.

Proof: Let $\epsilon>0$. We have

$$\begin{aligned}\frac{d}{d\epsilon}L\left(Q+\int_0^\epsilon\frac{D^*(Q_x,G)}{\dot{L}(Q_x)}dx\right)&=\dot{L}(Q_\epsilon)\frac{d}{d\epsilon}Q_\epsilon\\&=\dot{L}(Q_\epsilon)\frac{D^*(Q_\epsilon,G)}{\dot{L}(Q_\epsilon)}\\&=D^*(Q_\epsilon,G).\end{aligned}$$

This completest the proof of (11). The submodel statement was already shown above, but we now provided formal sufficient conditions. □

We refer to the appendix for an example demonstrating that the analytic formula (11) is indeed a submodel.

5 Example: One-step TMLE of effect among the treated

Let $O = (W, A, Y) \sim P_0$ and let $\mathcal{M}$ be a locally nonparametric statistical model. Let $\Psi : \mathcal{M} \to \mathbb{R}$ be defined by $\Psi(P) = E_P(E_P(Y \mid A = 1, W) - E_P(Y \mid A = 0, W) \mid A = 1)$. The efficient influence curve of Ψ at P is given by (Zheng et al., 2013):

$$D^*(P)(O) = H_1(g, q)(A, W)(Y - \bar{Q}(A, W)) + \frac{A}{q}\{\bar{Q}(1, W) - \bar{Q}(0, W) - \Psi(P)\},$$

where $g(a \mid W) = P(A = a \mid W)$, $\bar{Q}(a, W) = E_P(Y \mid A = a, W)$, $q = P(A = 1)$, and

$$H_1(g, q)(A, W) = \frac{A}{q} - \frac{(1 - A)g(1 \mid W)}{qg(0 \mid W)}.$$

We note that

$$\Psi(P) = \Psi_1(Q_W, \bar{Q}, g, q) = \int \{\bar{Q}(1, w) - \bar{Q}(0, w)\} \frac{g(1 \mid w)}{q} dQ_W(w),$$

where Q_W is the probability distribution of W under P. So, if we define $Q = (Q_W, \bar{Q}, g, q)$, then $\Psi(P) = \Psi_1(Q)$. For notational convenience, we will use $\Psi(P)$ and $\Psi(Q)$ interchangeably. Since we can estimate Q_W and q with their empirical probability distributions, we are only interested in a universal least favorable submodel for $(\bar{Q}, g)$. We can orthogonally decompose $D^*(P) = D_1^*(P) + D_2^*(P) + D_3^*(P)$ in $L_0^2(P)$ into scores of $\bar{Q}$, g, and Q_W, respectively, where

$$\begin{aligned}
D_1^*(P) &= \frac{g(1 \mid W)}{q}\{\bar{Q}(1, W) - \bar{Q}(0, W) - \Psi(Q)\} \\
D_2^*(P) &= H_2(Q)(W)(A - g(1 \mid W)) \\
D_3^*(P) &= \frac{g(1 \mid W)}{q}\{\bar{Q}(1, W) - \bar{Q}(0, W) - \Psi(Q)\},
\end{aligned}$$

and

$$H_2(Q)(W) = \frac{\bar{Q}(1, W) - \bar{Q}(0, W) - \Psi(Q)}{q}.$$

Thus the component of the efficient influence curve corresponding with $(\bar{Q}, g)$ is given by $D_1^*(Q) + D_2^*(Q)$.

We consider the following loss-functions and local least favorable submodels for $\bar{Q}$ and g (Zheng et al., 2013):

$$\begin{aligned}
L_1(\bar{Q})(O) &= -\{Y \log \bar{Q}(A,W) + (1-Y)\log(1-\bar{Q}(A,W))\} \\
\text{Logit}\bar{Q}_\epsilon^{lfm} &= \text{Logit}\bar{Q} - \epsilon H_1(g,q) \\
L_2(g)(O) &= -\{A \log g(1 \mid W) + (1-A)\log g(0 \mid W)\} \\
\text{Logit}\bar{g}_\epsilon^{lfm} &= \text{Logit}\bar{g} - \epsilon H_2(Q).
\end{aligned}$$

We now define the sum loss function $L(\bar{Q},g) = L_1(\bar{Q}) + L_2(g)$ and local least favorable submodel $\{Q_\epsilon^{lfm}, g_\epsilon^{lfm} : \epsilon\}$ through $(\bar{Q},g)$ at $\epsilon = 0$ satisfying

$$\frac{d}{d\epsilon} L(\bar{Q}_\epsilon^{lfm}, g_\epsilon^{lfm})\Big|_{\epsilon=0} = D_1^*(Q) + D_2^*(Q).$$

Thus, we can conclude that this defines indeed a local least favorable submodel for $(\bar{Q},g)$.

The universal least favorable submodel (3) is now defined by the following recursive definition: for $\epsilon \geq 0$ and $d\epsilon > 0$,

$$\begin{aligned}
\text{Logit}\bar{Q}_{\epsilon+d\epsilon} &= \text{Logit}\bar{Q}_{\epsilon,d\epsilon}^{lfm} \\
&= \text{Logit}\bar{Q}_\epsilon - d\epsilon H_1(g_\epsilon, q) \\
\text{Logit}\bar{g}_\epsilon &= \text{Logit}\bar{g}_{\epsilon,d\epsilon}^{lfm} \\
&= \text{Logit}\bar{g}_\epsilon - d\epsilon H_2(Q_W, \bar{Q}_\epsilon, q).
\end{aligned}$$

Similarly, we have a recursive relation for $\epsilon < 0$, but since all these formulas are just symmetric versions of the $\epsilon > 0$ case, we will focus on $\epsilon > 0$. This expresses the next $(Q_{\epsilon+d\epsilon}, g_{\epsilon+d\epsilon})$ in terms of previously calculated $(Q_x, g_x : x \leq \epsilon)$, thereby fully defining this universal least favorable submodel. This recursive definition corresponds with the following integral representation of this universal least favorable submodel:

$$\begin{aligned}
\text{Logit}\bar{Q}_\epsilon &= \text{Logit}\bar{Q} - \int_0^\epsilon H_1(g_x, q)dx \\
\text{Logit}\bar{g}_\epsilon &= \text{Logit}\bar{g} - \int_0^\epsilon H_2(Q_W, \bar{Q}_x, q)dx.
\end{aligned}$$

Let's now explicitly verify that this indeed satisfies the key property of a universal least favorable submodel. Clearly, it is a submodel and it contains (Q,g) at $\epsilon = 0$. The score of $\bar{Q}_\epsilon$ at ϵ is given by $H_1(g_\epsilon, q)(Y - \bar{Q}_\epsilon)$ and the score of g_ϵ at ϵ is given by $H_2(Q_W, \bar{Q}_\epsilon, q)(A - \bar{g}_\epsilon(W))$, so that

$$\begin{aligned}
\frac{d}{d\epsilon} L(\bar{Q}_\epsilon, g_\epsilon) &= H_1(g_\epsilon, q)(Y - \bar{Q}_\epsilon) + H_2(Q_W, \bar{Q}_\epsilon, q)(A - \bar{g}_\epsilon(W)) \\
&= D_1^*(Q_W, \bar{Q}_\epsilon, g_\epsilon, q) + D_2^*(Q_W, \bar{Q}_\epsilon, g_\epsilon, q),
\end{aligned}$$

explicitly proving that indeed this is a universal least favorable model for $(\bar{Q}, g)$.

In our previous work on the TMLE for the effect among the treated we implemented the TMLE based on the local least favorable submodel $\{\bar{Q}^{lfm}_{\epsilon_1}, \bar{g}^{lfm}_{\epsilon_2} : \epsilon_1, \epsilon_2\}$, using a separate ϵ_1 and ϵ_2 for $\bar{Q}$ and $\bar{g}$. This TMLE requires several iterations till convergence.

The TMLE based on the universal least favorable submodel above is implemented as follows, given an initial estimator $(\bar{Q}, g)$. One first determines the sign of the derivative at $h = 0$ of $P_n L(\bar{Q}_h, g_h)$. Suppose that the derivative is negative so that it decreases for $h > 0$. Then, one keeps iteratively calculating $(\bar{Q}_{\epsilon+d\epsilon}, g_{\epsilon+d\epsilon})$ for small $d\epsilon > 0$, given $(\bar{Q}_x, g_x : x \leq \epsilon)$, till $P_n L(\bar{Q}_{\epsilon+d\epsilon}, g_{\epsilon+d\epsilon}) \geq P_n L(\bar{Q}_\epsilon, g_\epsilon)$, at which point the desired local maximum ϵ_n is attained. The TMLE of $(\bar{Q}_0, g_0)$ is now given by $\bar{Q}_{\epsilon_n}, g_{\epsilon_n}$, which solves $P_n\{D^*_1(Q_{\epsilon_n}) + D^*_2(Q_{\epsilon_n})\} = 0$, where $Q_{\epsilon_n} = (Q_{W,n}, \bar{Q}_{\epsilon_n}, g_{\epsilon_n}, q_n)$, $Q_{W,n}, q_n$ are the empirical counterparts of $Q_{W,0}, q_0$. Since, we also have $P_n D^*_3(Q_{\epsilon_n}) = 0$, it follows that $P_n D^*(Q_{\epsilon_n}) = 0$. The (one-step) TMLE of $\Psi(Q_0)$ is given by the corresponding plug-in estimator $\Psi(Q_{\epsilon_n})$.

6 Universal canonical one-dimensional submodel that targets a multidimensional target parameter

Let $\Psi : \mathcal{M} \rightarrow H$ be a Hilbert-space value pathwise differentiable target parameter. Typically, we simply have $H = \mathbb{R}^d$ endowed with the standard inner product $\langle x, y \rangle = \sum_{j=1}^d x_j y_j$. However, we also allow that $\Psi(P)$ is a function $t \rightarrow \Psi(P)(t)$ from $\tau \subset \mathbb{R}$ to $\mathbb{R}$ in a Hilbert space $L^2(\Lambda)$ endowed with inner product $\langle h_1, h_2 \rangle = \int h_1(t) h_2(t) d\Lambda(t)$, where Λ is a user supplied positive measure with $\int d\Lambda(t) < \infty$. For notational convenience, we will often denote the inner product $\langle h_1, h_2 \rangle$ with $h_1^\top h_2$, analogue to the typical notation for the inner product in $\mathbb{R}^d$. Let $\parallel h \parallel = \sqrt{\langle h, h \rangle}$ be the Hilbert space norm, which would be the standard Euclidean norm in the case that $H = \mathbb{R}^d$. Let $D^*(P)$ be the canonical gradient. If $H = \mathbb{R}^d$, then this is a d-dimensioncal canonical gradient $D^*(P) = (D^*_j(P) : j = 1, \ldots, d)$, but in general $D^*(P) = (D^*_t(P) : t \in \tau)$. Let $L(P) = -\log p$, where $p = dP/d\mu$ is a density of $P \ll \mu$ w.r.t. some dominating measure μ. In this section we will construct a one-dimensional submodel $\{P_\epsilon : \epsilon \geq 0\}$ through P at $\epsilon = 0$ so that, for any $\epsilon \geq 0$,

$$\frac{d}{d\epsilon} P_n L(p_\epsilon) = \parallel P_n D^*(P_\epsilon) \parallel . \tag{6}$$

The one-step TMLE P_{ϵ_n} with $\epsilon_n = \arg\max_\epsilon P_n L(P_\epsilon)$, or ϵ_n chosen large enough so that the derivative is smaller than (e.g.) $1/n$, now solves $\frac{d}{d\epsilon_n} P_n L(P_{\epsilon_n}) = 0$ (or $< 1/n$), and thus $\parallel P_n D^*(P_{\epsilon_n}) \parallel = 0$ (or $< 1/n$). Note that $\parallel P_n D^*(P_{\epsilon_n}) \parallel = 0$ implies that $P_n D^*_t(P_{\epsilon_n}) = 0$ for all $t \in \tau$ so that the one-step TMLE solves all desired estimating equations.

6.1 A universal canonical submodel that targets a multidimensional target parameter

Consider the following submodel: for $\epsilon \geq 0$, we define

$$\begin{aligned} p_\epsilon &= p\Pi_{[0,\epsilon]}\left(1 + \frac{\{P_n D^*(P_x)\}^\top D^*(P_x)}{\| D^*(P_x) \|}dx\right) \\ &= p\exp\left(\int_0^\epsilon \frac{\{P_n D^*(P_x)\}^\top D^*(P_x)}{\| D^*(P_x) \|}dx\right). \end{aligned} \tag{7}$$

Similarly, for $\epsilon < 0$, we could define

$$\begin{aligned} p_\epsilon &= p\Pi_{[\epsilon,0]}\left(1 - \frac{\{P_n D^*(P_x)\}^\top D^*(P_x)}{\| D^*(P_x) \|}dx\right) \\ &= p\exp\left(-\int_\epsilon^0 \frac{\{P_n D^*(P_x)\}^\top D^*(P_x)}{\| D^*(P_x) \|}dx\right) \end{aligned}$$

but the only relevant direction for the TMLE is the one that moves in the direction in which the log-likelihood increases and that is the one for $\epsilon \geq 0$. Therefore, it suffices to only define the submodel for $\epsilon \geq 0$.

Theorem 4 *We have $\{p_\epsilon : \epsilon \geq 0\}$ is a family of probability densities, its score at ϵ is a linear combination of $D_t^*(P_\epsilon)$ for $t \in \tau$, and is thus in the tangent space at $T(P_\epsilon)$, and*

$$\frac{d}{d\epsilon}P_n L(P_\epsilon) = \| P_n D^*(P_\epsilon) \| .$$

As a consequence, we have $\frac{d}{d\epsilon}P_n L(P_\epsilon) = 0$ implies $\| P_n D^*(P_\epsilon) \| = 0$.

As before, our practical construction below demonstrates that, under regularity conditions, we actually have that $\{p_\epsilon : \epsilon\} \subset \mathcal{M}$ is also a submodel.

The normalization by $\| D^*(P_x) \|$ is motivated by a practical analogue construction below and provides an important intuition behind this analytic construction. However, we can replace this by any other normalization for which the derivative of the log-likelihood at ϵ equals a norm of $P_n D^*(P_\epsilon)$. To illustrate this let's consider the case that $H = \mathbb{R}^d$. For example, we could consider the following submodel. Let $\Sigma_n(P_x) = P_n\{D^*(P_x)D^*(P_x)^\top\}$ be the empirical covariance matrix of $D^*(P_x)$, and let $\Sigma_n^{-1}(P_x)$ be its inverse. We could then define for $\epsilon > 0$,

$$p_\epsilon = p\exp\left(\int_0^\epsilon \{P_n D^*(P_x)\}^\top \Sigma_n^{-1} D^*(P_x)dx\right),$$

In this case, we have

$$\frac{d}{d\epsilon}P_n L(P_\epsilon) = P_n D^*(P_\epsilon)^\top \Sigma_n(P_\epsilon)^{-1} P_n D^*(P_\epsilon).$$

This seems to be an appropriately normalized norm, equal to the euclidean norm of the orthonormalized version of the original $D^*(P_\epsilon)$, so that the one-step TMLE will still satisfy that $\parallel P_n D^*(P_{\epsilon_n}) \parallel = 0$.

It is not clear to us if these choices have a finite sample implication for the resulting one-step TMLE (asymptotics is the same), and if one choice would be better than another, but either way, the resulting one-step TMLE ends up with a P_{ϵ_n} satisfying $P_n D^*(P_{\epsilon_n}) = 0$ (or $o_P(1/\sqrt{n})$), the only key ingredient in the proof of the asymptotic efficiency of the TMLE.

6.2 The practical construction of a universal canonical one-dimensional submodel targeting a multidimensional target parameter

Let's define a local least favorable submodel $\{p_\delta^{lfm} : \delta\} \subset \mathcal{M}$ by the following local property:

$$\frac{d}{d\delta} \log p_\delta^{lfm} \Big|_{\delta=0}^{\top} \delta = D^*(P)^\top \delta.$$

For the case that $H = \mathbb{R}^d$, this corresponds with assuming that the score of the submodel at $\delta = 0$ equals the canonical gradient $D^*(P)$, while, for a general Hilbert space, it states that the derivative of $\log p_\epsilon$ in the direction δ (a function in H) equals $\langle D^*(P), \delta \rangle = \int D_t^*(P)\delta(t) d\Lambda(t)$.

Consider the log-likelihood criterion $P_n L(P_\delta^{lfm})$, and we note that its derivative at $\delta = 0$ in the direction δ equals $\langle P_n D^*(P), \delta \rangle = \{P_n D^*(P)\}^\top \delta$. For a small number dx, we want to maximize the log-likelihood over all δ with $\parallel \delta \parallel \leq dx$, and locally, this corresponds with maximizing its linear gradient approximation:

$$\delta \rightarrow (P_n D^*(P))^\top \delta.$$

By the Cauchy-Schwarz inequality, it follows that this is maximized over δ with $\parallel \delta \parallel \leq dx$ by

$$\delta_n^*(P, dx) = \frac{P_n D^*(P)}{\parallel P_n D^*(P) \parallel} dx \equiv \delta_n^*(P) dx,$$

where we defined $\delta_n^*(P) = P_n D^*(P) / \parallel P_n D^*(P) \parallel$. We can now define our update $P_{dx} = P_{\delta_n^*(P,dx)}^{lfm}$. This process can now be iterated by applying the above with P replaced by P_{dx}, resulting in an update P_{2dx}, and in general P_{Kdx}. So this updating process is defined by the differential equation:

$$P_{x+dx} = P_{x,\delta_n^*(P_x)dx)}^{lfm},$$

where $P_{x,\delta}^{lfm}$ is the local least favorable multidimensional submodel above but now through P_x instead of P.

Assuming that the local least favorable model $h \rightarrow p^{lfm}_{x,h}$ is continuously twice differentiable with a score $D^*(P_x)$ at $h=0$, we obtain a second order Tailor expansion

$$\begin{aligned} p^{lfm}_{x,\delta^*_n(P_x)dx} &= p_x + \left\{ \frac{d}{dh} p^{lfm}_{x,h} \Big|_{h=0} \right\}^\top \delta^*_n(P_x)dx + O((dx)^2) \\ &= p_x(1+\{\delta^*_n(P_x)\}^\top D^*(P_x)dx) + O((dx)^2), \end{aligned}$$

so that, under mild regularity conditions, we obtain

$$p_{x+dx} = p_x(1+\{\delta^*_n(P_x)\}^\top D^*(P_x)dx + O((dx)^2).$$

This implies:

$$p_x = p \exp\left(\int_0^\epsilon \frac{(P_n D^*(P_x))^\top}{\parallel P_n D^*(P_x) \parallel} D^*(P_x)dx \right).$$

So now we obtain the exact same representation (7) as above. Since the above practical construction starts out with $P \in \mathcal{M}$ and never leaves the model $\mathcal{M}$, this proves that, under mild regularity conditions, this analytic representation (2) is actually a submodel of $\mathcal{M}$ after all, but, when using its practical implementation and approximation, one should use the actual local least favorable submodel in order to guarantee that one stays in the model. We can formalize this in a theorem analogue to Theorem 2, but instead such a theorem will be presented in Section 8 for the more general targeted minimum loss-based estimation methodology.

The above practical construction provides us with an intuition for the normalization by $\parallel P_n D^*(P_x) \parallel$.

6.3 Existence of MLE or approximate MLE ϵ_n.

Since

$$P_n \log p_\epsilon = \int_0^\epsilon \parallel P_n D^*(P_x) \parallel dx,$$

and its derivative thus equals $\parallel P_n D^*(P_\epsilon) \parallel$, we have that the log-likelihood is non-decreasing in ϵ.

If the local least favorable submodel in the practical construction of the one-dimensional universal canonical submodel $\{p_\epsilon : \epsilon \geq 0\}$ (7) only contains densities with supremum norm smaller than some $M < \infty$ (e.g., this is assumed by the model $\mathcal{M}$), then we will have that $\sup_{\epsilon \geq 0} \sup_{o \in \mathcal{O}} p_\epsilon(o) < M < \infty$. This implies that $P_n \log p_\epsilon$ is bounded from above by $\log M$. Let's first assume that $\lim_{\epsilon \to \infty} P_n \log p_\epsilon < \infty$. Thus, the log-likelihood is an strictly increasing function till it becomes flat, if ever. Suppose that $\limsup_{x\to\infty} \parallel P_n D^*(P_x) \parallel > \delta > 0$ for some $\delta > 0$. Then it follows that the log-likelihood converges to infinity when ϵ converges to infinity, which contradicts the assumption that the log-likelihood is bounded from above by $\log M < \infty$. Thus, we know that $\limsup_{x\to\infty} \parallel P_n D^*(P_x) \parallel = 0$ so that we can find an ϵ_n so that for $\epsilon > \epsilon_n \parallel P_n D^*(P_\epsilon) \parallel < 1/n$, as desired.

Suppose now that we are in a case in which the log-likelihood converges to infinity when $\epsilon \to \infty$, so that our bounded log likelihood assumption is violated. This might correspond with a case in which each p_ϵ is a continuous density, but p_ϵ starts approximating an empirical distribution when $\epsilon \to \infty$. Even in such a case, one would expect that we will have that $\parallel P_n D^*(P_\epsilon) \parallel \to 0$, just like an NPMLE of a continuous density of a survival time solves the efficient influence curve equation for its survival function.

The above practical construction of the submodel, as an iterative local maximization of the log-likelihood along its gradient, strongly suggests that even without the above boundedness assumption the derivative $\parallel P_n D^*(P_\epsilon) \parallel$ will converge to zero as $\epsilon \to \infty$ so that the desired MLE or approximate MLE exists.

6.4 A universal score-specific one-dimensional submodel targeting a multivariate score equation

In the above two subsections we could simply replace $D^*(P)$ by a user supplied $D(P)$, giving us a theoretical one-dimensional parametric model $\{P_\epsilon : \epsilon\}$ so that the derivative $\frac{d}{d\epsilon} P_n L(P_\epsilon)$ at ϵ equals $\parallel P_n D(P_\epsilon) \parallel$, so that a corresponding one-step TMLE will solve $P_n D(P_{\epsilon_n}) = 0$. Similarly, given a local parametric model whose score at $\epsilon = 0$ equals $D(P)$ will yield a corresponding practical construction of this universal submodel. One can also use such a universal score-specific submodel to construct one-step TMLE of a one-dimensional target parameter with extra properties by making it solve not only the efficient influence curve equation but also other equations of interest (such as the $P_n D_2$ in Section C). A variety of such types of TMLE have been proposed in the literature using an iterative TMLE based on a local score-specific submodel.

7 Example: A one-step TMLE, based on universal canonical one-dimensional submodel, of an infinite dimensional target parameter

An open problem has been the construction of an efficient substitution estimator $\Psi(P_n^*)$ of a pathwise differentiable infinite dimensional target parameter $\Psi(P_0)$ such as a survival function. Current approaches would correspond with incompatible estimators such as using a TMLE for each $\Psi_t(P_0)$ separately, resulting in a non-substitution estimator such as a non-monotone estimator of a survival function. In this section we demonstrate, through a causal inference example, that our universal canonical submodel allows us to solve this problem with a one-step TMLE.

Let $O = (W, A, T) \sim P_0$, where W are baseline covariates, $A \in \{0, 1\}$ is a point-treatment, and T is a survival time. Consider a statistical model $\mathcal{M}$ that only makes assumptions about the conditional distribution $g_0(a \mid W) = P_0(A = a \mid W)$ of A, given W. Let $W \to d(W) \in \{0, 1\}$ be a given dynamic treatment satisfying

$g_0(d(W) \mid W) > 0$ a.e. Let $\Psi : \mathcal{M} \to H$ be defined by:

$$\Psi(P)(t) = E_P P(T > t \mid A = d(W), W),\ t \geq 0.$$

Under a causal model and the randomization assumption this equals the counterfactual survival function $P(T_d > t)$ of the counterfactual survival time T_d under intervention d.

Let H be the Hilbert space of real valued functions on $\mathbb{R}_{\geq 0}$ endowed with inner product $h_1^\top h_2 = \langle h_1, h_2 \rangle = \int h_1(t) h_2(t) d\Lambda(t)$ for some user-supplied positive and finite measure Λ. The norm on this Hilbert space is thus given by $\parallel h \parallel = \sqrt{hh^\top} = \sqrt{\int h(t)^2 d\Lambda(t)}$. Let $\bar{Q}_t(A, W) = P(T > t \mid A, W)$, $Y(t) = I(T > t)$, and Q_W the marginal distribution of W. The efficient influence curve $D^*(P) = (D_t^*(P) : t \geq 0)$ is defined by:

$$\begin{aligned} D_t^*(P)(O) &= \frac{I(A = d(W))}{g(A \mid W)}(Y(t) - \bar{Q}_t(A, W)) + \{\bar{Q}_t(d(W), W) - \Psi(P)(t)\} \\ &\equiv D_{1,t}^*(g, \bar{Q}) + D_{2,t}^*(P), \end{aligned}$$

where $D_{1,t}^*(g, \bar{Q})$ is the first component of the efficient influence curve that is a score of the conditional distribution of T, given A, W. Notice that $\Psi(P) = \Psi_1(Q_W, \bar{Q}) = (Q_W \bar{Q}_t : t \geq 0)$. We will estimate $Q_{W,0}$ with the empirical distribution of $W_1, \ldots, W_n$, so that a TMLE will only need to target the estimator of the conditional survival function $\bar{Q}_0$ of T, given A, W. Let $q(t \mid A, W)$ be the density of T, given A, W and let q_n be an initial estimator of this conditional density. For example, one might use machine learning to estimate the conditional hazard $q_0/\bar{Q}_0$, which then implies a corresponding density estimator q_n. We are also given an estimator g_n of g_0.

The universal canonical one-dimensional submodel (7) applied to $p = q_n$ is defined by the following recursive relation: for $\epsilon > 0$,

$$q_{n,\epsilon} = q_n \exp\left(\int_0^\epsilon \frac{\{P_n D_1^*(g_n, \bar{Q}_{n,x})\}^\top D_1^*(g_n, \bar{Q}_{n,x})}{\parallel D_1^*(g_n, \bar{Q}_{n,x}) \parallel} dx \right).$$

To get some more insight in this expression, we note, for example, that the inner product is given by:

$$\{P_n D_1^*(g_n, \bar{Q}_{n,x})\}^\top D_1^*(g_n, \bar{Q}_{n,x})(o) = \int_t (P_n D_{1,t}^*(g_n, \bar{Q}_{n,x}) D_{1,t}^*(g_n, \bar{Q}_{n,x})(o) d\Lambda(t), \tag{8}$$

and similarly we have such an integral representation of the norm in the denominator. Our theorem 4, or explicit verification, shows that for all $\epsilon \geq 0$ $q_{n,\epsilon}$ is a conditional density of T, given A, W, and

$$\frac{d}{d\epsilon} P_n \log q_{n,\epsilon} = \parallel P_n D_1^*(g_n, \bar{Q}_{n,\epsilon}) \parallel.$$

Thus, if we move ϵ away from zero, the log-likelihood increases, and, one searches for the first ϵ_n so that this derivative is smaller than (e.g.) $1/n$. Let $q_n^* = q_{n,\epsilon_n}$, and let $\bar{Q}_{n,t}^*(A, W) = (\int_t^\infty q_n^*(s \mid A, W)ds : t \geq 0)$ be its corresponding conditional survival function. Then our one-step TMLE of the d-specific survival function $\Psi(P_0)$ is given by $\psi_n^* = \Psi(Q_{W,n}, \bar{Q}_n^*) = Q_{W,n}\bar{Q}_n^*$:

$$\psi_n^*(t) = \frac{1}{n}\sum_{i=1}^{n} \bar{Q}_{n,t}^*(d(W_i), W_i).$$

Since q_n^* is an actual conditional density, it follows that ψ_n^* is a survival function. Suppose, that the derivative of the log-likelihood at ϵ_n equals zero exactly (instead of $1/n$). Then, we have $\| P_n D^*(g_n, Q_{W,n}, \bar{Q}_n^*) \| = 0$, so that for each $t \geq 0$, $P_n D_t^*(g_n, Q_{W,n}, \bar{Q}_n^*) = 0$, making $\psi_n^*(t)$ a standard TMLE of $\psi_0(t)$, so that its asymptotic linearity can be established accordingly, and that proof can be easily extended to be uniformly in all $t \in \tau$. In this manner, under the previously mentioned regularity conditions and a second order term condition (now uniformly in t), we obtain

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

where, by assumption, $\sup_t \mid R_n(t) \mid = o_P(1/\sqrt{n})$. In particular, if g_0 is known, then the second order term condition is automatically satisfied. This asymptotic linearity proves the asymptotic efficiency of the substitution estimator ψ_n^* as an estimator of the infinite dimensional survival function. The asymptotic linearity (and its required Donsker class condition) implies that $\sqrt{n}(\psi_n^* - \psi_0)$ converges weakly to a Gaussian process with covariance function $\rho(s, t) = P_0 D_s^*(P_0) D_t^*(P_0)$. This also allows the construction of a simultaneous confidence band for ψ_0. Due to the double robustness of the efficient influence curve, one can also obtain asymptotic linearity with an inefficient influence curve under misspecfication of either g_n or $\bar{Q}_n$.

If we only have $\| P_n D^*(P_n^*) \| = 1/n$ (instead of 0), then the above proof still applies but with a second order term R_n for which now $\| R_n \| = o_P(1/\sqrt{n})$, so that we obtain asymptotic efficiency in the Hilbert space norm, beyond the pointwise efficiency of $\psi_n^*(t)$. However, in practice, one can actually track the supremum norm $\| P_n D^*(P_{\epsilon_n}) \|_\infty = \sup_t \mid P_n D_t^*(P_{\epsilon_n}) \mid$, and if one observes that for the selected ϵ_n this supremum norm is smaller than $1/n$, then, we still obtain the asymptotic efficiency in supremum norm above.

Regarding the practical construction of $q_{n,\epsilon}$, we could use the following infinite dimensional local least favaroble submodel through a conditional density q given by

$$q_\delta^{lfm} = q(1 + \delta^\top D_1^*(g, \bar{Q})),$$

and follow the practical construction described in the previous section for general local least favorable submodels. Notice that here $\delta^\top D_1^*(g, \bar{Q}) = \int \delta(t) D_{1,t}^*(g, \bar{Q}) d\Lambda(t)$. In order to guarantee that the supremum norm of the density q_δ^{lfm} for local δ with

$\| \delta \| < dx$ remains below a universal constant $M < \infty$, one could present such models in the conditional hazard on a logistic scale that bounds the hazard between $[0, M]$. However, we doubt that this will be an issue in practice, and may be it is necessary that the continuous density $q_{n,\epsilon}$ approximates an empirical distribution is some sense in order to solve $\| P_n D^*(P_\epsilon) \| = 0$, in which case we do not want to obstruct this to happen.

8 Universal canonical one-dimensional submodel for targeted minimum loss-based estimation of a multidimensional target parameter

8.1 A universal canonical one-dimensional submodel

For the sake of presentation we will focus on the case that the target parameter is Euclidean values, i.e. $H = \mathbb{R}^d$, but the presentation immediately generalizes to infinite dimensional target parameters, as in the previous section. Let's now generalize this construction to a universal canonical submodel for the more general targeted minimum loss based estimation methodology. We now assume that $\Psi(P) = \Psi_1(Q(P)) \in \mathbb{R}^d$ for some target parameter $Q : \mathcal{M} \to Q(\mathcal{M})$ defined on the model and real valued function $\Psi_1 : Q(\mathcal{M}) \to \mathbb{R}^d$. Let $Q(\mathcal{M}) = \{Q(P) : P \in \mathcal{M}\}$ be the parameter space of this parameter. Let $L(Q)(O)$ be a loss-function for $Q(P)$ in the sense that $Q(P) = \arg\min_{Q \in Q(\mathcal{M})} PL(Q)$. Let $D^*(P) = D^*(Q(P), G(P))$ be the canonical gradient of Ψ at P, where $G : \mathcal{M} \to G(\mathcal{M})$ is some nuisance parameter. We consider the case that the linear span of the components of the efficient influence curve $D^*(P)$ is in the tangent space of Q, so that a least favorable submodel does not need to fluctuate G: otherwise, one just includes G in the definition of Q. Given, (Q, G), let $\{Q_\delta^{lfm} : \delta\} \subset Q(\mathcal{M})$ be a local d-dimensional least favorable model w.r.t. loss function $L(Q)$ at $\delta = 0$ so that

$$\left. \frac{d}{d\delta} L(Q_\delta^{lfm}) \right|_{\delta=0} = D^*(Q, G).$$

The dependence of this submodel on G is suppressed in this notation.

Consider the empirical risk $P_n L(Q_\delta^{lfm})$, and we note that its gradient at $\delta = 0$ equals $P_n D^*(Q, G)$. For a small number dx, we want to maximize the empirical risk over all δ with $\| \delta \| \leq dx$, and locally, this corresponds with maximizing its linear gradient approximation:

$$\delta \to (P_n D^*(Q, G))^\top \delta.$$

By the Cauchy-Schwarz inequality, it follows that this is maximized over δ with $\| \delta \| \leq dx$ by

$$\delta_n^*(Q, dx) = \frac{P_n D^*(Q, G)}{\| P_n D^*(Q, G) \|} dx \equiv \delta_n^*(Q) dx,$$

where we defined $\delta_n^*(Q) = P_n D^*(Q,G)/ \parallel P_n D^*(Q,G) \parallel$. We can now define our update $Q_{dx} = Q^{lfm}_{\delta_n^*(Q,dx)}$. This process can now be iterated by applying the above with Q replaced by Q_{dx}, resulting in an update Q_{2dx}, and in general Q_{Kdx}. So this updating process is defined by the differential equation:

$$Q_{x+dx} = Q^{lfm}_{x,\delta_n^*(Q_x)dx)},$$

where $Q^{lfm}_{x,\delta}$ is the local least favorable multidimensional submodel above but now through Q_x instead of Q.

Assume that for some $\dot{L}(Q)(O)$, we have

$$\frac{d}{dh} L(Q^{lfm}_{x,h})\Big|_{h=0} = \dot{L}(Q_x)\ \frac{d}{dh} Q^{lfm}_{x,h}\Big|_{h=0}. \tag{9}$$

Then,

$$\frac{d}{dh} Q^{lfm}_{x,h}\Big|_{h=0} = \frac{D^*(Q_x,G)}{\dot{L}(Q_x)}.$$

Utilizing that the local least favorable model $h \rightarrow Q^{lfm}_{x,h}$ is continuously twice differentiable with a score $D^*(Q_x,G)$ at $h=0$, we obtain a second order Tailor expansion

$$\begin{aligned} Q^{lfm}_{x,\delta_n^*(Q_x)dx} &= Q_x + \frac{d}{dh} Q^{lfm}_{x,h}\Big|_{h=0} \delta_n^*(Q_x)dx + O((dx)^2) \\ &= Q_x + \frac{D^*(Q_x,G)^\top}{\dot{L}(Q_x)} \delta_n^*(Q_x)dx + O((dx)^2). \end{aligned}$$

This implies the following recursive analytic definition of the universal canonical submodel through Q:

$$Q_\epsilon = Q + \int_0^\epsilon \frac{D^*(Q_x,G)^\top}{\dot{L}(Q_x)} \delta_n^*(Q_x)dx. \tag{10}$$

Let's now explicitly verify that this indeed satisfies the desired condition so that the one-step TMLE solves $P_n D^*(Q_{\epsilon_n}, G) = 0$. Only assuming (9) it follows that

$$\begin{aligned}
\frac{d}{d\epsilon} P_n L(Q_\epsilon) &= P_n \frac{d}{d\epsilon} L(Q_\epsilon) \\
&= P_n \dot{L}(Q_\epsilon) \frac{d}{d\epsilon} Q_\epsilon \\
&= P_n \dot{L}(Q_\epsilon) \frac{D^*(Q_\epsilon, G)^\top}{\dot{L}(Q_\epsilon)} \delta_n^*(Q_\epsilon) \\
&= P_n D^*(Q_\epsilon, G)^\top \delta_n^*(Q_\epsilon) \\
&= \{P_n D^*(Q_\epsilon, G)\}^\top \frac{P_n D^*(Q_\epsilon, G)}{\parallel P_n D^*(Q_\epsilon, G) \parallel} \\
&= \frac{\sum_{j=1}^d \{P_n D_j^*(Q_\epsilon, G)\}^2}{\parallel P_n D^*(Q_\epsilon, G) \parallel} \\
&= \parallel P_n D^*(Q_\epsilon, G) \parallel .
\end{aligned}$$

In addition, under some regularity conditions, so that the following derivation in terms of the local least favorable submodel applies, it also follows that $Q_\epsilon \in Q(\mathcal{M})$.

This proves the following theorem.

Theorem 5 *Given, any (Q, G) compatible with model $\mathcal{M}$, let $\{Q_\delta^{lfm} : \delta \in B_a(0)\} \subset Q(\mathcal{M})$ be a local least favorable model w.r.t. loss function $L(Q)$ at $\delta = 0$ so that*

$$\left. \frac{d}{d\delta} L(Q_\delta^{lfm}) \right|_{\delta=0} = D^*(Q, G).$$

Here $B_a(0) = \{x : \parallel x \parallel < a\}$ for some positive number a. Assume that for some $\dot{L}(Q)(O)$, we have

$$\left. \frac{d}{d\epsilon} L(Q_\epsilon^{lfm}) \right|_{\epsilon=0} = \dot{L}(Q) \left. \frac{d}{d\epsilon} Q_\epsilon^{lfm} \right|_{\epsilon=0} .$$

Consider the corresponding univariate model $\{Q_\epsilon : \epsilon\}$ defined by (10). It goes through Q at $\epsilon = 0$, and, it satisfies that for all ϵ

$$P_n \frac{d}{d\epsilon} L(Q_\epsilon) = \parallel P_n D^*(Q_\epsilon, G) \parallel, \tag{11}$$

where $\parallel x \parallel = \sqrt{\sum_{j=1}^d x_j^2}$ is the Euclidean norm.

In addition, assume that a in $B_a(0)$ can be chosen to be independent of (Q, G) in $\{(Q_\epsilon, G) : \epsilon > 0\}$, and assume the following second order Tailor expansion: for $h = (h_1, \ldots, h_d)$,

$$\begin{aligned}
Q_{\epsilon,h}^{lfm} &= Q_\epsilon + \left. \frac{d}{dh} Q_{\epsilon,h}^{lfm} \right|_{h=0} h + R_2(Q_\epsilon, G, \parallel h \parallel) \\
&= Q_\epsilon + \frac{D^*(Q_\epsilon, G)}{\dot{L}(Q_\epsilon)} h + R_2(Q_\epsilon, G, \parallel h \parallel),
\end{aligned}$$

where

$$\sup_{\epsilon}\sup_{o\in\mathcal{O}} \mid R_2(Q_\epsilon, G, \parallel h \parallel)(o) \mid = O((\parallel h \parallel^2).$$

We also assume that $\sup_\epsilon \sup_{o\in\mathcal{O}} \frac{|D^*(Q_\epsilon,G)}{\dot{L}(Q_\epsilon)}(o) \mid < \infty$.

Then, we also have $\{Q_\epsilon : \epsilon \geq 0\} \subset \mathcal{M}$.

9 Universal canonical one-dimensional submodel for targeted minimum loss-based estimation of a multidimensional target parameter when the loss function depends on nuisance parameters

9.1 A universal canonical one-dimensional submodel

Let's now generalize this construction of a universal canonical submodel in the previous section to a parameter Q whose loss-function depends on a nuisance parameter. As in the previous section we assume that $\Psi(P) = \Psi_1(Q(P)) \in \mathbb{R}^d$ for some target parameter $Q : \mathcal{M} \to Q(\mathcal{M})$ defined on the model and real valued function $\Psi_1 : Q(\mathcal{M}) \to \mathbb{R}^d$. Let $L_{\Gamma P}(Q)(O)$ be a loss-function for $Q(P)$ in the sense that $Q(P) = \arg\min_{Q\in Q(\mathcal{M})} PL_{\Gamma(P)}(Q)$, where $\Gamma : \mathcal{M} \to \Gamma(\mathcal{M})$ is some nuisance parameter. Let $D^*(P) = D^*(Q(P), G(P))$ be the canonical gradient of Ψ at P, where $G : \mathcal{M} \to G(\mathcal{M})$ is some nuisance parameter. We consider the case that the linear span of the components of the efficient influence curve $D^*(P)$ is in the tangent space of Q, so that a least favorable submodel does not need to fluctuate G: otherwise, one just includes G in the definition of Q. One will have that $\Gamma(P)$ only depends on P through $(Q(P), G(P))$, so that we will also use the notation $\Gamma(Q, G)$. Given, (Q, G), let $\{Q^{lfm}_\delta : \epsilon\} \subset Q(\mathcal{M})$ be a local d dimensional least favorable model w.r.t. loss function $L_{\Gamma(Q,G)}(Q)$ at $\delta = 0$ so that

$$\frac{d}{d\delta} L_{\Gamma(Q,G)}(Q^{lfm}_\delta)\bigg|_{\epsilon=0} = D^*(Q, G).$$

The dependence of this submodel on G is suppressed in this notation.

Consider the empirical risk $P_n L_{\Gamma(Q,G)}(Q^{lfm}_\delta)$, and we note that its gradient at $\delta = 0$ equals $P_n D^*(Q, G)$. For a small number dx, we want to maximize the empirical risk over all δ with $\parallel \delta \parallel \leq dx$, and locally, this corresponds with maximizing its linear gradient approximation:

$$\delta \to (P_n D^*(Q, G))^\top \delta.$$

By the Cauchy-Schwarz inequality, it follows that this is maximized over δ with $\parallel \delta \parallel \leq dx$ by

$$\delta^*_n(Q, dx) = \frac{P_n D^*(Q, G)}{\parallel P_n D^*(Q, G) \parallel} dx \equiv \delta^*_n(Q) dx,$$

where we defined $\delta_n^*(Q) = P_n D^*(Q,G)/ \parallel P_n D^*(Q,G) \parallel$. We can now define our update $Q_{dx} = Q^{lfm}_{\delta_n^*(Q,dx)}$. This process can now be iterated by applying the above with Q replaced by Q_{dx} and $\Gamma(Q,G)$ replaced by $\Gamma(Q_{dx},G)$, resulting in an update Q_{2dx}, and in general Q_{Kdx}. So at the k-th step, we have

$$Q_{x+kdx} = Q^{lfm}_{x+(k-1)dx,\delta_n^*(Q_{(k-1)dx}))dx},$$

where

$$\delta_n^*(Q_{(k-1)dx}) = P_n D^*(Q_{(k-1)dx},G)/ \parallel P_n D^*(Q_{(k-1)dx},G) \parallel .$$

So this updating process is defined by the differential equation:

$$Q_{x+dx} = Q^{lfm}_{x,\delta_n^*(Q_x)dx)},$$

where $Q^{lfm}_{x,\epsilon}$ is the local least favorable multidimensional submodel above but now through Q_x instead of Q.

Assume that for some $\dot{L}_\Gamma(Q)(O)$, we have

$$\frac{d}{dh} L_{\Gamma_x}(Q^{lfm}_{x,h})\Big|_{h=0} = \dot{L}_{\Gamma_x}(Q_x) \ \frac{d}{dh} Q^{lfm}_{x,h}\Big|_{h=0}, \tag{12}$$

where we used the notation $\Gamma_x = \Gamma(Q_x,G)$. Then,

$$\frac{d}{dh} Q^{lfm}_{x,h}\Big|_{h=0}) = \frac{D^*(Q_x,G)}{\dot{L}_{\Gamma_x}(Q_x)}.$$

Utilizing that the local least favorable model $h \rightarrow Q^{lfm}_{x,h}$ is continuously twice differentiable with a score $D^*(Q_x,G)$ at $h=0$, we obtain a second order Tailor expansion

$$\begin{aligned} Q^{lfm}_{x,\delta_n^*(Q_x)dx} &= Q_x + \frac{d}{dh} Q^{lfm}_{x,h}\Big|_{h=0} \delta_n^*(Q_x)dx + O((dx)^2) \\ &= Q_x + \frac{D^*(Q_x,G)^\top}{\dot{L}_{\Gamma_x}(Q_x)} \delta_n^*(Q_x)dx + O((dx)^2). \end{aligned}$$

This implies the following recursive analytic definition of the universal least favorable model through Q:

$$Q_\epsilon = Q + \int_0^\epsilon \frac{D^*(Q_x,G)^\top}{\dot{L}_{\Gamma_x}(Q_x)} \delta_n^*(Q_x)dx. \tag{13}$$

Let's now explicitly verify that this submodel defined by (13) indeed satisfies the desired condition that the one-step TMLE Q_{ϵ_n} with ϵ_n defined as the value closest to zero for which

$$\frac{d}{dh} P_n L_{\Gamma(Q_\epsilon,G)}(Q_{\epsilon+h})\Big|_{h=0} = 0$$

solves $\| P_n D^*(Q_{\epsilon_n}, G) \| = 0$. Only assuming (12) it follows that

$$\begin{aligned}
\frac{d}{dh} P_n L_{\Gamma_\epsilon}(Q_{\epsilon+h}) \Big|_{h=0} &= P_n \frac{d}{dh} L_{\Gamma_\epsilon}(Q_{\epsilon+h}) \Big|_{h=0} \\
&= P_n \dot{L}_{\Gamma_\epsilon}(Q_\epsilon) \frac{d}{d\epsilon} Q_\epsilon \\
&= P_n \dot{L}_{\Gamma_\epsilon}(Q_\epsilon) \frac{D^*(Q_\epsilon, G)^\top}{\dot{L}_{\Gamma_\epsilon}(Q_\epsilon)} \delta_n^*(Q_\epsilon) \\
&= P_n D^*(Q_\epsilon, G)^\top \delta_n^*(Q_\epsilon) \\
&= \{P_n D^*(Q_\epsilon, G)\}^\top \frac{P_n D^*(Q_\epsilon, G)}{\| P_n D^*(Q_\epsilon, G) \|} \\
&= \| P_n D^*(Q_\epsilon, G) \| .
\end{aligned}$$

So this proves that indeed this submodel and the corresponding one-step TMLE (which updates the loss through Γ_ϵ when moving along ϵ) indeed solves $\| P_n D^*(Q_{\epsilon_n}, G) \| = 0$.

In addition, under some regularity conditions, so that the above derivation in terms of the local least favorable submodel applies, it also follows that $Q_\epsilon \in Q(\mathcal{M})$. This proves the analogue of Theorem 5.

9.2 Example: One-step TMLE of parameters of marginal structural working model for multiple time-point interventions

In this subsection we develop a new one-step TMLE based on the universal canonical one-dimensional submodel, while the previous closed form TMLE developed in(Petersen et al., 2013) was based on a local least favorable submodel with d-parameters at each time point.

Suppose that the observed data structure is $O = (L(0), A(0), L(1), A(1), Y) \sim P_0$, where $Y \in \{0,1\}$ or continuous with $Y \in (0,1)$. Let $V = f(L(0))$ be some potential baseline effect modifier of interest. Suppose that our statistical model $\mathcal{M}$ only makes assumptions about $g_0 = (g_{0,A(0)}, g_{0,A(1)})$. Consider a set of dynamic treatment regimens $\mathcal{D}$, and for a $d \in \mathcal{D}$, let $E_0(Y_d \mid V)$ be the conditional mean of Y_d, given V, under the G-computation formula $p_0^d = q_{L(0)} q_{L(1)} q_Y d_{A(0)} d_{A(1)}$ obtained by replacing $g_{0,A(0)}, g_{0,A(1)}$ in the factorization of the density p_0 of P_0 by the degenerate conditional distributions $d_{A(0)}$ and $d_{A(1)}$. Here $Q_{L(0)}$ is the marginal distribution of $L(0)$, and $Q_{L(1)}, Q_Y$ are the conditional densities of $L(1)$, given $A(0), L(0)$, and of Y, given $\bar{L}(1), \bar{A}(1)$, respectively, while $q_{L(0)}, q_{L(1)}, q_Y$ are their respective densities. Given a working model $\{m_\beta : \beta \in \mathbb{R}^d\}$ for $E_0(Y_d \mid V)$, and weight function $(d, V) \to h(d, V)$, the target parameter $\Psi : \mathcal{M} \to \mathbb{R}^d$ is defined by

$$\Psi(P) = \arg\min_\psi E_P \sum_{d \in \mathcal{D}} h(d, V) L^F(m_\psi(d, V))(Y_d, V),$$

where $L^F(m)(Y_d, V) = -\{Y_d \log m(d,V) + (1-Y_d)\log(1-m(d,V))\}$ is the log-likelihood loss function for $E(Y_d \mid V)$. By the sequential regression representation of $E_P(Y_d \mid V)$ (Bang and Robins, 2005), it follows that $\Psi(P) = \Psi_1(Q_{L(0)}, \bar{Q})$, where $\bar{Q} = (\bar{Q}^1, \bar{Q}^2) = (\bar{Q}^{1,d}, \bar{Q}^{2,d} : d \in \mathcal{D})$, and

$$\begin{aligned}\bar{Q}^{d,1}(\bar{L}(1)) &= E_P(Y \mid \bar{L}(1), \bar{A}(1) = \bar{d}(\bar{L}(1)))\\ \bar{Q}^{d,0}(L(0)) &= E_P(\bar{Q}^{d,1}(\bar{L}(1)) \mid L(0), A(0) = d_{A(0)}(L(0))).\end{aligned}$$

We assume that $\text{Logit} m_\beta(d,V) = \beta^\top \phi(d,V)$ for some vector of basis functions $\phi = (\phi_1, \ldots, \phi_d)$. The efficient influence curve of Ψ at P is given by $D^*(Q,G) = c(\Psi(Q))^{-1} D(Q,G)$, where

$$\begin{aligned}D(Q,G)(O) &= \sum_{d\in\mathcal{D}} h_1(d,V)(\bar{Q}^{d,1}(d,L(0)) - m_{\Psi(Q)}(d,V))\\ &+ \sum_{d\in\mathcal{D}} h_1(d,V)\frac{I(A(0) = d_{A(0)}(L(0))}{g_{A(0)}(O)}(\bar{Q}^{d,1}(\bar{L}(1)) - \bar{Q}^{d,0}(L(0)))\\ &+ \sum_{d\in\mathcal{D}} h_1(d,V)\frac{I(\bar{A}(1) = \bar{d}(\bar{L}(1)))}{g_{A(0)}g_{A(1)}(O)}(Y - \bar{Q}^{d,1}(\bar{L}(1)))\\ &\equiv D^0(Q) + D^1(\bar{Q},G) + D^2(\bar{Q},G),\end{aligned}$$

and

$$\begin{aligned}h_1(d,V) &= h(d,V)\phi(d,V)\\ c(\psi) &= E_P \sum_{d\in\mathcal{D}} h(d,V)\phi(d,V)\phi(d,V)^\top m_\psi(1-m_\psi)(d,V).\end{aligned}$$

Consider the following loss functions for the components of $\bar{Q} = (\bar{Q}^1, \bar{Q}^2) = (\bar{Q}^{1,d}, \bar{Q}^{2,d} : d \in \mathcal{D})$:

$$\begin{aligned}L_2(\bar{Q}^2) &= -\sum_{d\in\mathcal{D}} I(\bar{A}(1) = \bar{d}(\bar{L}(1)))\{Y \log \bar{Q}^{2,d} + (1-Y)\log(1-\bar{Q}^{2,d})\}\\ L_{1,\bar{Q}^2}(\bar{Q}^1) &= -\sum_{d\in\mathcal{D}} I(A(0) = d_0(L(0)))\{\bar{Q}^{2,d}\log \bar{Q}^{1,d} + (1-\bar{Q}^{2,d})\log(1-\bar{Q}^{1,d})\}.\end{aligned}$$

Given $\bar{Q}^1_n$, we will estimate $\bar{Q}^0_0$ with $\bar{Q}^0_n = P_n \bar{Q}^1_n$, an empirical mean. As a consequence, we only need a TMLE of $\bar{Q}^2_0$ and $\bar{Q}^1_0$, and the TMLE of $\bar{Q}^0_0$ follows by taking the empirical mean over $L(0)$ of the TMLE of $\bar{Q}^1_0$.

We can now define the sum loss function for $\bar{Q}$:

$$L_{\bar{Q}^{2,\prime}}(\bar{Q}) \equiv L_2(\bar{Q}^2) + L_{1,\bar{Q}^{2\prime}}(\bar{Q}^1),$$

which is indexed by nuisance parameter $\bar{Q}^2$ itself. For notational convenience, let's denote this nuisance parameter with $\Gamma(\bar{Q}) = \bar{Q}^2$. Then, this loss-function can also be represented as:

$$L_\gamma(\bar{Q}^2, \bar{Q}^1) = L_2(\bar{Q}^2) + L_{1,\gamma}(\bar{Q}^1).$$

Indeed, we have $L_{\gamma_0}(\bar{Q})$ is a valid loss function for $\bar{Q}_0 = \arg\min_{\bar{Q}} P_0 L_{\gamma_0}(\bar{Q})$.

Consider the following local least favorable d-dimensional submodel through $\bar{Q} = (\bar{Q}^2, \bar{Q}^1)$:

$$\begin{aligned}\text{Logit}\bar{Q}_\delta^{2,d,lfm} &= \text{Logit}\bar{Q}^{2,d} - \delta^\top H_2(d,g)\\ \text{Logit}\bar{Q}_\delta^{1,d,lfm} &= \text{Logit}\bar{Q}^{1,d} - \delta^\top H_1(d,g)\end{aligned}$$

where $H_2(d,g) = h_1(d,V)I(\bar{A}(1) = \bar{d}(\bar{L}(1)))/(g_{A(0)}g_{A(1)}(O))$, and $H_1(d,g) = h_1(d,V)I(A(0) = d_0(L(0)))/g_{A(0)}(O)$. Indeed, we have

$$\frac{d}{d\delta} L_{\bar{Q}^2}(\bar{Q}_\delta^{lfm})\Big|_{\delta=0} = \bar{D}(\bar{Q}, G) \equiv D^1(\bar{Q}, G) + D^2(\bar{Q}, G).$$

Let dx be given. Define the d-dimensional vector

$$\delta_n^*(\bar{Q}) = \frac{P_n \bar{D}(Q,G)}{\parallel P_n \bar{D}(Q,G) \parallel}.$$

We can now define our first update $\bar{Q}_{dx} = \bar{Q}^{lfm}_{\delta_n^*(\bar{Q})dx}$. In other words, for each $d \in \mathcal{D}$, we have

$$\begin{aligned}\text{Logit}\bar{Q}_{dx}^{2,d} &= \text{Logit}\bar{Q}^{2,d} - \delta_n^*(\bar{Q}) dx H_2(d,g)\\ \text{Logit}\bar{Q}_{dx}^{1,d} &= \text{Logit}\bar{Q}^{1,d} - \delta_n^*(\bar{Q}) dx H_1(d,g).\end{aligned}$$

We can now iterate this updating process. So let

$$\delta_n^*(\bar{Q}_{dx}) = \frac{P_n \bar{D}(\bar{Q}_{dx}, G)}{\parallel P_n \bar{D}(Q_{dx}, G) \parallel}$$

We can now define our second update $\bar{Q}_{2dx} = \bar{Q}^{lfm}_{dx, \delta_n^*(\bar{Q}_{dx})dx}$. In other words, for each $d \in \mathcal{D}$, we have

$$\begin{aligned}\text{Logit}\bar{Q}_{2dx}^{2,d} &= \text{Logit}\bar{Q}_{dx}^{2,d} - \delta_n^*(\bar{Q}_{dx}) dx H_2(d,g)\\ &= \text{Logit}\bar{Q}^{2,d} - \delta_n^*(\bar{Q}) dx H_2(d,g) - \delta_n^*(\bar{Q}_{dx}) dx H_2(d,g)\\ &= \text{Logit}\bar{Q}^{d,d} - \sum_{k=0}^{1} \delta_n^*(\bar{Q}_{kdx}) dx H_2(d,g)\\ \text{Logit}\bar{Q}_{2dx}^{1,d} &= \text{Logit}\bar{Q}^{1,d} - \sum_{k=0}^{1} \delta_n^*(\bar{Q}_{kdx})) dx H_1(d,g).\end{aligned}$$

So, by iteration it follows that the desired universal one-dimensional submodel is given by

$$\bar{Q}_\epsilon = \bar{Q}^{lfm}_{\int_0^\epsilon \delta_n^*(\bar{Q}_x)dx}.$$

Let's define the d-dimensional vector

$$C_n(\epsilon) = \int_0^\epsilon \frac{P_n \bar{D}(\bar{Q}_x, G)}{\parallel P_n \bar{D}(\bar{Q}_x, G) \parallel} dx.$$

Then the desired universal canonical one-dimensional submodel can be presented as follows: for each $d \in \mathcal{D}$, and $\epsilon > 0$,

$$\begin{aligned} \text{Logit} \bar{Q}^{2,d}_\epsilon &= \text{Logit} \bar{Q}^{2,d} - C_n(\epsilon)^\top H_2(d, g) \\ \text{Logit} \bar{Q}^{1,d}_\epsilon &= \text{Logit} \bar{Q}^{1,d} - C_n(\epsilon)^\top H_1(d, g). \end{aligned}$$

Let's now explicitly verify that the one-step TMLE indeed solves $P_n \bar{D}(\bar{Q}_{\epsilon_n}, G) = 0$ at $\epsilon_n > 0$ defined by the smallest $\epsilon > 0$ for which $\frac{d}{dh} P_n L_{\bar{Q}^2_\epsilon}(\bar{Q}_{\epsilon+h}) \Big|_{h=0} = 0$. Here we use that the empirical risk decreases in ϵ. Let

$$C_n'(\epsilon) = \frac{d}{d\epsilon} C_n(\epsilon) = \frac{P_n \bar{D}(\bar{Q}_\epsilon, G)}{\parallel P_n \bar{D}(\bar{Q}_\epsilon, G) \parallel}.$$

We have

$$\begin{aligned} \frac{d}{dh} P_n L_{\bar{Q}^2_\epsilon}(\bar{Q}_{\epsilon+h}) \Big|_{h=0} &= P_n \sum_{d \in \mathcal{D}} h_1(d, V) C_n'(\epsilon)^\top H_1(d, g)(\bar{Q}^{1,d} - \bar{Q}^{0,d}) \\ &\quad + P_n \sum_{d \in \mathcal{D}} h_1(d, V) C_n'(\epsilon)^\top H_2(d, g)(Y - \bar{Q}^{2,d}) \\ &= \frac{P_n \bar{D}(\bar{Q}_\epsilon, G)^\top}{\parallel P_n \bar{D}(\bar{Q}_\epsilon, G) \parallel} P_n \bar{D}(\bar{Q}_\epsilon, G) \\ &= \parallel P_n \bar{D}(\bar{Q}_\epsilon, G) \parallel . \end{aligned}$$

This proves that it is indeed a submodel that satisfies the desired condition so that the TMLE of $\Psi(P_0)$ is given by the one-step TMLE $\Psi_1(Q_{L(0),n}, \bar{Q}_{\epsilon_n})$.

10 Concluding remark

Given a d-variate estimating function $(Q, O) \rightarrow D(Q, G)(O)$, a loss function $L(Q)$ for $Q : \mathcal{M} \rightarrow Q(\mathcal{M})$, a local d-dimensional submodel $\{Q^{sm}_\delta : \delta\} \subset Q(\mathcal{M})$ so that $\frac{d}{d\delta} L(Q^{sm}_\delta)\big|_{\delta=0} = D(Q, G)$, we constructed a one-dimensional universal submodel $\{Q_\epsilon : \epsilon \geq 0\} \subset Q(\mathcal{M})$ through Q, at $\epsilon = 0$, that has the property that for all $\epsilon \geq 0$ $\frac{d}{d\epsilon} P_n L(Q_\epsilon) = \parallel P_n D(Q_\epsilon, G) \parallel$, where $\parallel \cdot \parallel$ is the Euclidean norm. Our analytic formula for this universal submodel does not depend on the local submodel, but the local submodel can still play a role for the practical construction. In the special case $d = 1$, we also constructed a universal one-dimensional submodel so that for all ϵ $\frac{d}{d\epsilon} L(Q_\epsilon) = D(Q_\epsilon, G)$, which then implies $\frac{d}{d\epsilon} P_n L(Q_\epsilon) = P_n D(Q_\epsilon, G)$. For each of these universal submodels, the one-step TMLE Q_{ϵ_n} with $\epsilon_n = \arg\min_\epsilon P_n L(Q_\epsilon)$

solves each $P_n D_j(Q_{\epsilon_n}, G) = 0$, $j = 1, \ldots, d$. We showed how this result immediately extends to an infinite dimensional estimating function $D = (D_t : t \in \tau)$, by replacing the Euclidean inner product by an Hilbert space inner product. If $D()$ is the canonical gradient of a target parameter, we referred to this submodel as the universal canonical submodel, and, if $d = 1$, the universal least favorable submodel.

The constructions of these universal submodels correspond with iteratively defining $Q_{\epsilon+d\epsilon} = Q^{sm}_{\epsilon,\delta(\epsilon)d\epsilon}$ where $\delta(\epsilon) = P_n D(Q_\epsilon, G)/ \parallel P_n D(Q_\epsilon, G) \parallel$ moves along the gradient of the empirical risk $P_n L(Q_\epsilon)$ at ϵ. These practical constructions demonstrate that this algorithm succeeds in updating an initial Q into an update $Q_n^* = Q_{\epsilon_n}$ that solves the desired equation $P_n D(Q_{\epsilon_n}, G) = 0$ while *minimally* decreasing the empirical risk relative to its initial value $P_n L(Q)$. That is, with minimal additional data fitting it achieves the desired goal, while fully preserving the statistical properties of the initial estimator represented by Q.

The universal submodels have dramatic implications for the TMLE literature by allowing one to construct one-step TMLE for any multivariate and even infinite dimensional pathwise differentiable target parameters, solving the desired estimating equation, so that this TMLE is asymptotically efficient and possibly has additional desired properties implied by solving the equation $P_n D(Q_{\epsilon_n}, G) = 0$. The one-step TMLE step only involves minimizing an empirical risk over a univariate fluctuation parameter ϵ. In the current literature, we proposed defined various iterative TMLE based on multivariate local submodels that can now be replaced by a more stable one-step TMLE only relying on maximizing over a univariate ϵ. We demonstrated such new one-step TMLE for various examples in this article, but obviously this will impact many more problems than the ones presented here.

The important advantages of the TMLE based on a local least favorable submodel relative to estimating equation methods and the one-step estimator have been emphasized in the literature. Since the estimating equation methodology is more limited than the one-step estimator by 1) relying on an estimating function representation of the efficient influence curve, 2) existence and 3) uniqueness of its solution, let's focus on contrasting the TMLE to the one-step estimator. One important advantage of the TMLE relative to the one-step estimator has been that it is a substitution estimator thereby making it in principle more robust by respecting the global constraints of the model $\mathcal{M}$. Beyond this, the fact that the TMLE updates an initial estimator through minimization of a loss-function specific empirical risk, it allows one to further refine the targeted update step such as carried out in C-TMLE. Another advantage is that it actually provides a corresponding data distribution $P_n^* \in \mathcal{M}$ compatible with the estimator of the target parameter, for example, allowing one to compare different TMLE by the empirical risk of P_n^*. On the other hand, the one-step estimator takes only one step, and that can add important stability relative to a possibly iterative TMLE, making the comparison not so clear in the case that the TMLE is iterative. However, our new universal submodels presented in this article make the TMLE also a single-step estimator, thereby dealing with this possible criticism of TMLE.

The benefit of being a substitution estimator is particularly appealing if one estimates an infinite dimensional target parameter such as a survival function with clear global structure. Due to our universal canonical one-dimensional submodel, we could provide one-step TMLE that complete respects this global structure of the infinite dimensional target parameter, something a one-step estimator (or estimating equation method) can not achieve.

Future simulation studies will have to evaluate the practical benefits that come with the new one-step TMLEs based on universal least favorable or canonical submodels, relative to TMLEs based on the typical local least favorable submodel.

References

H. Bang and J.M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61:962–972, 2005.

P.J. Bickel, C.A.J. Klaassen, Y. Ritov, and J. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag, 1997.

M. Carone, I. Diaz, and Mark J. van der Laan. Higher-order targeted minimum loss-based estimation. Technical Report 331, www.bepress.com/ucbbiostat/paper331, University of California, Berkeley, 2014.

I. Diaz, M. Carone, and Mark J. van der Laan. Second order inference for the mean of a variable missing at random. Technical Report 337, www.bepress.com/ucbbiostat/paper337, University of California, Berkeley, 2015.

Iván Díaz and Mark J van der Laan. Targeted data adaptive estimation of the causal dose response curve. *Journal of Causal Inference*, 2013, In press.

S. Gruber and M.J. van der Laan. An application of collaborative targeted maximum likelihood estimation in causal inference and genomics. *Int J Biostat*, 6(1), 2010.

S. Gruber and M.J. van der Laan. Targeted minimum loss based estimator that outperforms a given estimator. *The International Journal of Biostatistics*, 8(1): Article 11, doi: 10.1515/1557–4679.1332, 2012.

Samuel D Lendle, Bruce Fireman, and Mark J van der Laan. Balancing score adjusted targeted minimum loss-based estimation. 2013.

M. Petersen, J. Schwab, S. Gruber, N. Blaser, M. Schomaker, and M.J. van der Laan. Targeted minimum loss based estimation of marginal structural working models. *Journal of Causal Inference*, submitted, technical report http://biostats.bepress.com/ucbbiostat/paper312/, 2013.

E.C. Polley, Sherri Rose, and M.J. van der Laan. Super learning. In M.J. van der Laan and S. Rose, editors, *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York Dordrecht Heidelberg London, 2012.

J.M. Robins and A. Rotnitzky. Recovery of information and adjustment for dependent censoring using surrogate markers. In *AIDS Epidemiology*, Methodological issues. Bikhäuser, 1992.

S. Rose and M.J. van der Laan. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York, 2011.

D.B. Rubin and M.J. van der Laan. Empirical efficiency maximization: Improved locally efficient covariate adjustment in randomized experiments and survival analysis. *The International Journal of Biostatistics*, Vol. 4, Iss. 1, Article 5, 2008.

R.J.C.M. Starmans. Models, inference and truth: Probabilistic reasoning in the information era. In M.J. van der Laan and S. Rose, editors, *Targeted Learning: Causal Inference for Observational and Experimental Studies*, pages 1–20. Springer, New York, 2011.

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.

M.J. van der Laan. Estimation based on case-control designs with known prevalance probability. *The International Journal of Biostatistics*, page http://www.bepress.com/ijb/vol4/iss1/17/, 2008.

M.J. van der Laan. Statistical inference when using data adaptive estimators of nuisance parameters. Technical Report 302, Division of Biostatistics, University of California, Berkeley, submitted to IJB, 2012a.

M.J. van der Laan. Causal inference for networks. Technical Report 300, Division of Biostatistics, University of California, Berkeley, 2012b. Submitted to JCI, also technical report http://www.bepress.com/ucbbiostat/paper300.

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*, 2010.

M.J. van der Laan and M.L. Petersen. Targeted learning. In *Ensemble Machine Learning*, chapter pages 117–156, ISBN 978-1-4419-9326-7. Springer, New York, 2012.

M.J. van der Laan and J.M. Robins. *Unified methods for censored longitudinal data and causality.* Springer, New York, 2003.

M.J. van der Laan and D. Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006.

M.J. van der Laan, S. Dudoit, and A.W. van der Vaart. The cross-validated adaptive epsilon-net estimator. *Statistics and Decisions*, 24(3):373–395, 2006.

M.J. van der Laan, E. Polley, and A. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(25), 2007. ISSN 1.

A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Emprical Processes.* Springer-Verlag New York, 1996.

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.

W. Zheng and M.J. van der Laan. Cross-validated targeted minimum loss based estimation. In M.J. van der Laan and S. Rose, editors, *Targeted Learning: Causal Inference for Observational and Experimental Studies.* Springer, New York, 2011.

W. Zheng, M.L. Petersen, and M.J. van der Laan. Estimating the effect of a community-based intervention with two communities. *Journal of Causal Inference*, 1(Issue 1):83–106, 2013.

Appendix

A Example for Section 3: Universal least favorable submodel for parametric models, and resulting one-step TMLE

This section represents the final subsection of 3.

Even though the standard MLE for a parametric model is asymptotically efficient for any pathwise differentiable target parameter, if the dimension of the finite dimensional parameter is high relative to sample size, then the MLE is often not well defined or overly variable so that regularization is needed, and in that case a TMLE is still needed. High dimensional linear regression is an example of such types of high dimensional parametric models, but also saturated models when O is discrete (but possibly with many possible values). This type of application of TMLE motivates us to consider the universal least favorable submodel and corresponding one-step TMLE for parametric models.

Let $O \sim P_{\theta_0} \in \mathcal{M} = \{P_\theta : \theta \in \Theta \subset \mathbb{R}^d\}$ be modeled with a d-dimensional parametric model. Assume that the model is dominated by a single dominating measure μ, and the density $dP_\theta/d\mu$ will be denoted with p_θ. Let $\Psi : \mathcal{M} \rightarrow \mathbb{R}$ be a real valued target parameter, which is pathwise differentiable with canonical gradient $D^*(P_\theta)$ at $P_\theta \in \mathcal{M}$. Let $S_j(P_\theta) = \frac{d}{d\theta_j} \log dP_\theta/d\mu$ be the score of θ_j, $j = 1, \ldots, d$. The tangent space $T(P_\theta)$ at P_θ is the linear span of these d scores. Let $\alpha(P_\theta) = (\alpha_j(P_\theta) : j = 1, \ldots, d)$ be a uniquely defined vector of scalars such that

$$D^*(P_\theta) = \sum_{j=1}^{d} \alpha_j(P_\theta) S_j(P_\theta).$$

Such a vector $\alpha(P_\theta)$ exist and is unique if the $d \times d$ information matrix $I(P_\theta) = P_\theta S_\theta S_\theta^\top$ is invertible, but even when the tangent space is of lower dimension than d, there exist a whole space of such vectors of scalars, and this just selects one of them in a unique manner.

A local least favorable model $\{P_{\theta,\epsilon}^{lfm} : \epsilon\}$ through P_θ at $\epsilon = 0$ is given by:

$$P_{\theta,\epsilon}^{lfm} = P_{\theta+\epsilon\alpha(P_\theta)} = P_{(\theta_j+\epsilon\alpha_j(P_\theta):j=1,\ldots,d)}.$$

Let

$$\theta^{lfm}(\epsilon) = \theta + \epsilon\alpha(P_\theta)$$

be the corresponding least favorable path in the Θ space, so that we can denote $P_{\theta,\epsilon}^{lfm} = P_{\theta^{lfm}(\epsilon)}$. Indeed,

$$\begin{aligned}
\frac{d}{d\epsilon} \log p_{\theta,\epsilon}^{lfm}\Big|_{\epsilon=0} &= \frac{1}{p_\theta} \frac{d}{d\epsilon} p_{\theta+\epsilon\alpha(P_\theta)}\Big|_{\epsilon=0} \\
&= \sum_{j=1}^{d} \frac{1}{p_\theta} \frac{d}{d\theta_j} p_\theta \, \frac{d}{d\epsilon} (\theta_j + \epsilon\alpha_j(P_\theta))\Big|_{\epsilon=0} \\
&= \sum_{j=1}^{d} \alpha_j(P_\theta) S_j(P_\theta) \\
&= D^*(P_\theta).
\end{aligned}$$

Let the universal least favorable model through θ be defined by the following differential equation: for $\epsilon > 0, d\epsilon > 0$

$$\theta(\epsilon + d\epsilon) = \theta(\epsilon)^{lfm}(d\epsilon) = \theta(\epsilon) + d\epsilon\alpha(P_{\theta(\epsilon)}).$$

Similarly, we define $\theta(\epsilon - d\epsilon)$ for $\epsilon < 0$. The corresponding integral equation is given by: for $\epsilon > 0$ we have

$$\theta(\epsilon) = \theta + \int_0^\epsilon \alpha(P_{\theta(x)}) dx.$$

This differential or integral equation allows one to solve recursively for $\theta(\epsilon)$, given previous values $\theta(x)$ for $x < \epsilon$.

A corresponding universal least favorable submodel $\{P_{\theta,\epsilon} : \epsilon\}$ through P_θ is now defined by: for $\epsilon \geq 0$

$$\begin{aligned} P_{\theta,\epsilon} &= P_{\theta(\epsilon)} \\ &= P_{\theta+\int_0^\epsilon \alpha(P_{\theta(x)})dx}. \end{aligned}$$

And similarly we can define $P_{\theta,\epsilon}$ for $\epsilon < 0$. By our results, we also know that we could define this universal least favorable submodel through P_θ by: for $\epsilon \geq 0$

$$P_{\theta,\epsilon} = P_\theta \exp\left(\int_0^\epsilon D^*(P_{\theta,x})dx\right),$$

but for the sake of practical approximation one should prefer the above formulation in terms of a local least favorable submodel.

Suppose that $P_n \log p_{\theta_n^{lfm}(\epsilon)}$ is decreasing at $\epsilon = 0$. Then, the TMLE is defined by an initial estimator θ_n, and defining ϵ_n as smallest local maximum larger than 0 of $\epsilon \rightarrow P_n \log p_{\theta_n(\epsilon)}$. The TMLE of θ_0 is now given by $\theta_n^* = \theta_n(\epsilon_n)$, and the TMLE of $\Psi(P_{\theta_0})$ is given by $\Psi(P_{\theta_n^*})$.

B Example for Section 4 demonstrating that analytic formula (11) for universal least favorable submodel is indeed a submodel

Suppose $O = (W, A, Y) \sim P_0$, $A \in \{0, 1\}$ binary, Y binary in $\{0, 1\}$ or $Y \in (0, 1)$, and let the statistical model $\mathcal{M}$ be the nonparametric model or any model that only restricts the tangent space of the conditional distribution of A, given W. Let $\Psi : \mathcal{M} \rightarrow \mathbb{R}$ be defined by $\Psi(P) = E_P E_P(Y \mid A = 1, W)$. The efficient influence curve $D^*(P)(O) = A/\bar{g}(W)(Y - \bar{Q}(W)) + \bar{Q}(W) - \Psi(P)$, where $\bar{g}(W) = P(A = 1|W)$ and $\bar{Q}(W) = E_P(Y \mid A = 1, W)$. We note that $\Psi(P) = \Psi_1(Q) = Q_W \bar{Q}$, where $Q = (Q_W, \bar{Q})$, and Q_W is the probability distribution of W under P. We can decompose $D^*(P) = D_1^*(\bar{Q}, \bar{g}) + D_0^*(Q)$, where $D_1^*(\bar{Q}, \bar{g}) = A/\bar{g}(Y - \bar{Q}(W))$ is a score of the conditional distribution of Y, given A, W, while $D_0^*(Q)$ is a score of the marginal distribution of W. Since we estimate $Q_{W,0}$ with the empirical probability distribution of $W_1, \ldots, W_n$, there is no need to construct a submodel through Q_W, so that we focus on constructing a submodel through $\bar{Q}$ only.

A valid loss function for $\bar{Q}$ is given by

$$L(\bar{Q})(O) = -I(A = 1)\{Y \log \bar{Q}(W) + (1 - Y) \log(1 - \bar{Q}(W))\}.$$

Consider the local least favorable submodel through $\bar{Q}$:

$$\text{Logit}\bar{Q}_\epsilon^{lfm} = \text{Logit}\bar{Q} - \epsilon H(\bar{g}),$$

where $H(\bar{g})(A,W) = A/\bar{g}(W)$. This is indeed a local least favorable submodel for $\bar{Q}$ since

$$\frac{d}{d\epsilon} L(\bar{Q}_\epsilon^{lfm}) \Big|_{\epsilon=0} = D_1^*(\bar{Q}, \bar{g}).$$

Let's now compute the corresponding theoretical universal least favorable submodel (4). We have

$$\frac{d}{d\epsilon} L(\bar{Q}_\epsilon) = \frac{d}{d\epsilon} Q_\epsilon \left\{ -I(A=1) \frac{Y - \bar{Q}_\epsilon}{\bar{Q}_\epsilon (1 - \bar{Q}_\epsilon)} \right\}.$$

Thus,

$$\dot{L}(Q_\epsilon) = -I(A=1) \frac{Y - \bar{Q}_\epsilon}{\bar{Q}_\epsilon (1 - \bar{Q}_\epsilon)}.$$

Thus, the universal least favorable submodel (4) through Q is given by:

$$\bar{Q}_\epsilon = \bar{Q} - H(\bar{g}) \int_0^\epsilon \bar{Q}_x (1 - \bar{Q}_x) dx.$$

This integral equation shows that

$$\frac{\frac{d}{d\epsilon} \bar{Q}_\epsilon}{\bar{Q}_\epsilon (1 - \bar{Q}_\epsilon)} = -H(\bar{g}).$$

This has as solution $\bar{Q}_\epsilon = Q_\epsilon^{lfm}$, and since there is only one solution, this proves that the universal least favorable submodel $\bar{Q}_\epsilon = Q_\epsilon^{lfm}$. Indeed, it follows directly that for all ϵ

$$\frac{d}{d\epsilon} L(\bar{Q}_\epsilon^{lfm}) = D_1^*(\bar{Q}_\epsilon^{lfm}, \bar{g}),$$

showing that our local least favorable submodel is already a universal least favorable submodel. Indeed, the TMLE using Q_ϵ^{lfm} requires only one step. In particular, as predicted by our theory, this demonstrates that the analytic formula (4) respects the constraints that $\bar{Q} \in (0,1)$, even though that is not immediately obvious from its analytic integral or differential representation.

C Universal score-specific submodel generalizing the universal least favorable submodel

This section could be read after Section 5

Consider the above setting $O \sim P_0 \in \mathcal{M}$, $\Psi : \mathcal{M} \rightarrow \mathbb{R}$, $\Psi(P) = \Psi_1(Q(P))$, $Q(P) = \arg\min_Q PL(Q)$, Ψ is pathwise differentiable at P with canonical gradient $D^*(Q(P), G(P))$ for some nuisance parameter G that is orthogonal to Ψ in the sense that the nuisance tangent space of G is orthogonal to the tangent space of Q.

In Section 5 constructed universal least favorable models $\{Q_\epsilon : \epsilon\}$ for any loss-based parameter Q whose loss-based score $\frac{d}{d\epsilon} L(Q_\epsilon)$ at ϵ equals the efficient influence curve $D^*(Q_\epsilon, G)$. Using this universal least favorable submodel through an initial estimator of Q_0 results in a TMLE that takes only one step, and, as any TMLE, is asymptotically efficient under regularity conditions.

Let $L_2(G)$ be a loss function for G so that $G(P) = \arg\min_{G_1 \in G(\mathcal{M})} PL_2(G_1)$. Let $L(Q,G) = L(Q) + L_2(G)$ be the sum loss-function for (Q,G). Let $D_2(Q,G)$ be a user supplied element of the tangent space $T_G(P)$ of G in $L^2_0(P)$. Let's define a local score-specific (i.e., $D_2()$-specific) submodel $\{G^{sm}_\epsilon : \epsilon\} \subset G(\mathcal{M})$ as a submodel through G at $\epsilon = 0$ satisfying

$$\frac{d}{d\epsilon} L_2(G_\epsilon)\Big|_{\epsilon=0} = D_2(Q,G).$$

Then, given a local least favorable submodel $\{Q^{lfm}_\epsilon : \epsilon\}$ through Q, we have that $\{(Q^{lfm}_\epsilon, G^{sm}_\epsilon) : \epsilon\} \subset (Q,G)(\mathcal{M})$ satisfies

$$\frac{d}{d\epsilon} L(Q^{lfm}_\epsilon, G^{sm}_\epsilon)\Big|_{\epsilon=0} = D(Q,G) \equiv D^*(Q,G) + D_2(Q,G).$$

We refer to such a submodel $\{Q^{lfm}_\epsilon, G^{sm}_\epsilon : \epsilon\}$ as a local $D()$-specific submodel.

Typically, Q can be decomposed as $Q = (Q_1, Q_2)$ in which Q_{20} can always be consistently estimated, and one select $D_2(Q = (Q_1,Q_2), G)$ so that $D_2(Q_1, Q_{20}, G_0)$ equals minus the projection of $D^*(Q_1, Q_{20}, G_0)$ onto a subspace of the tangent space of G in $L^2(P_0)$. Such a choice implies that 1) for any Q_1 $D^*(Q_1, Q_{20}, G_0) - D_2(Q_1, Q_{20}, G_0)$ is a desired influence curve with significantly smaller variance than $D^*(Q_1, Q_{20}, G_0)$ at misspecified Q_1 and 2) $D^*(Q_{10}, Q_{20}, G_0) + D_2(Q_{10}, Q_{20}, G_0) = D^*(Q_{10}, Q_{20}, G_0)$. That is, D_2 yields a correction to a misspecified $D^*(Q_1, Q_{20}, G_0)$ that only kicks in when Q_1 is misspecified. In this way, the model is still a local least favorable submodel so that the TMLE is asymptotically efficient when both Q_0, G_0 are consistently estimated.

Specifically, one might be given a user supplied influence curve $D^0(Q_1, Q_{20}, G_0)$ at P_0 (for any given Q_1), which one can represent as

$$D^0(Q_1, Q_{20}, G_0) = D^*(Q_1, Q_{20}, G_0) + D_2(Q_1, Q_{20}, G_0),$$

for some $D_2(Q_1, Q_{20}, G_0) \in T_G(P_0)$. One can now define the desired score as:

$$\begin{aligned} D(Q_1, Q_{20}, G_0) &= D^*(Q_1, Q_{20}, G_0) \\ &\quad - \frac{P_0\{D^*(Q_1, Q_{20}, G_0) D_2(Q_1, Q_{20}, G_0)\}}{P_0 D_2^2(Q_1, Q_{20}, G_0)} D_2(Q_1, Q_{20}, G_0). \end{aligned}$$

This influence curve has smaller or equal variance than $D^0(Q_1, Q_{20}, G_0)$, and if $Q = Q_0$ (i.e. $Q_1 = Q_{10}$), then $D(Q_1, Q_{20}, G_0) = D^*(Q_0, G_0)$ is the efficient influence curve. By using this as the desired score equation, one will obtain a one-step TMLE that will be more efficient than an estimator with the user supplied influence curve $D^0(Q_1, Q_{20}, G_0)$ at P_0.

Such a TMLE is analyzed by using that $P_n D(Q_{1n}^*, Q_{2n}, G_n^*) = 0$,

$$P_0 D(Q_{1n}^*, Q_{2n}, G_n^*) = \Psi(Q_0) - \Psi(Q_n^*) + R_{2n},$$

for a second order term in $(Q_{2n} - Q_{20})$ and $G_n^* - G_0$, even when Q_{1n}^* is inconsistent for Q_{10}, so that

$$\Psi(Q_n^*) - \Psi(Q_0) = (P_n - P_0) D(Q_{1n}^*, Q_{2n}, G_n^*) + R_{2n}.$$

If now $R_{2n} = o_P(1/\sqrt{n})$, $D(Q_{1n}^*, Q_{2n}, G_n^*)$ falls in a P_0-Donsker class, $P_0\{D(Q_{1n}^*, Q_{2n}, G_n^*) - D(Q_1, Q_{20}, G_0)\}^2 \to 0$ in probability, then it follows that

$$\Psi(Q_n^*) - \Psi(Q_0) = (P_n - P_0) D(Q_1, Q_{20}, G_0) + o_P(1/\sqrt{n}).$$

In particular, if $Q_1 = Q_{10}$ it is asymptotically efficient, but even at misspecified Q_1 it has a desired influence curve $D(Q_1, Q_{20}, G_0)$.

In the current literature such TMLE have always been iterative TMLE, using more fitting than needed for the desired asymptotic properties(Gruber and van der Laan, 2012; Lendle et al., 2013). This motivates us again to define a universal score-specific (i.e., $D()$-specific) submodel as a submodel $\{(Q_\epsilon, G_\epsilon) : \epsilon\} \subset (Q, G)(\mathcal{M})$ so that for all ϵ

$$\frac{d}{d\epsilon} L(Q_\epsilon, G_\epsilon) = D^*(Q_\epsilon, G_\epsilon) + D_2(Q_\epsilon, G_\epsilon).$$

Such a universal submodel is defined by the recursive differential equation definition:

$$(Q_{\epsilon+d\epsilon}, G_{\epsilon+d\epsilon}) = (Q_{\epsilon,d\epsilon}^{lfm}, G_{\epsilon,d\epsilon}^{sm}),$$

where we need to keep in mind that the submodel $Q_{\epsilon,h}^{lfm}$ uses G_ϵ it its definition (if it depends on G), and, similarly, the submodel $G_{\epsilon,h}^{sm}$ uses Q_ϵ in its definition. As in the previous sections, this can be used to generate an analytic integral representation. However, in most applications such integral representations follow immediately, so that we just present the above recursive differentiable equation relation. Since

$$\left.\frac{d}{d\delta} L(Q_{\epsilon,\delta}^{lfm}, G_{\epsilon,\delta}^{sm})\right|_{\delta=0} = D^*(Q_\epsilon, G_\epsilon) + D_2(Q_\epsilon, G_\epsilon),$$

it follows that this submodel is indeed a universal score-specific submodel.

As before a TMLE using this universal score-specific submodel for updating (Q, G) will only require one step, and the TMLE $(Q_{\epsilon_n}, G_{\epsilon_n})$ will solve the desired score equation

$$0 = P_n D(Q_{\epsilon_n}, G_{\epsilon_n}) = P_n\{D^*(Q_{\epsilon_n}, G_{\epsilon_n}) + D_2(Q_{\epsilon_n}, G_{\epsilon_n})\},$$

so that it can be analyzed as above showing that, under regularity conditions, it is asymptotically linear with influence curve $D(Q_1, Q_{20}, G_0)$, which equals the efficient influence curve if Q_1 happens to be the true value Q_{10}.

C.1 Example: Targeting the treatment mechanism in TMLE for the additive treatment effect to obtain a more efficient estimator at misspecified Q

Let $O = (W, A, Y) \sim P_0$ and let $\mathcal{M}$ be a model that puts at most restrictions on the conditional probability distribution $g_0(a \mid W) = P_0(A = a \mid W)$. Let $\Psi : \mathcal{M} \rightarrow \mathbb{R}$ be defined by $\Psi(P) = E_P\{E_P(Y \mid A = 1, W) - E_P(Y \mid A = 0, W)\}$. We have that $\Psi(P) = \Psi_1(Q) = \Psi_1(Q_W, \bar{Q})$ is only a function of the distribution Q_W of W and the conditional mean $\bar{Q}$ of Y, given A, W. Let $D^*(Q, g)$ be the efficient influence curve at P, and let $D^*_a(\bar{Q}, g) = H_1(g)(Y - \bar{Q})$ be the corresponding efficient score for $\bar{Q}$, while $D^*_b(Q) = \bar{Q}(1, W) - \bar{Q}(0, W) - \Psi(Q)$ is the corresponding efficient score of Q_W, so that $D^*(Q, g) = D^*_a(\bar{Q}, g) + D^*_b(Q)$. For a given $\bar{Q}$, let

$$D_2(\bar{Q}, Q_{20}, g_0) = -\Pi(D^*_a(\bar{Q}, g_0) \mid T_2(g_0)),$$

where $T_2(g_0) \subset L^2_0(P_0)$ is a subspace of the nonparametric tangent space of g at P_0 consisting of all functions of (A, W) with conditional mean zero, given W, and Π denotes the projection operator onto $T_2(g_0)$ in the Hilbert space $L^2_0(P_0)$. Since a function of W is orthogonal to a function of A, W that has mean zero, given W, we also have that

$$\begin{aligned} D_2(\bar{Q}, Q_{20}, g_0) &= -\Pi(D^*_a(\bar{Q}, g_0) + D^*_b(Q_{W0}, \bar{Q}) \mid T_2(g_0)) \\ &= -\Pi(D^*(Q_{W0}, \bar{Q}, g_0) \mid T_2(g_0)). \end{aligned}$$

For example, $T_2(g_0)$ could be the tangent space of a parametric model through g_0. In the latter case this projection depends on covariances under P_0 so that Q_{20} indicates this dependence on P_0 beyond g_0, and it is clear that Q_{20} can be consistently estimated. We assume that $T_2(g_0)$ (i.e., $\mathcal{G}$) is small enough so that the projection operator (i.e. Q_{20}) can indeed be consistently estimated. This projection can be represented as $D_2(\bar{Q}, Q_{20}, g_0)(A, W) = H_2(\bar{Q}, Q_{20}, g_0)(W)(A - \bar{g}_0(W))$ for some H_2. The TMLE will now be tailored to solve $P_n\{D^*(Q_{W,n}, \bar{Q}^*_n, g^*_n) + D_2(\bar{Q}^*_n, Q_{2n}, g^*_n)\} = 0$, where $Q_{W,n}$ is the unbiased empirical distribution of $Q_{W,0}$, Q_{2n} is the unbiased estimator of the covariances coded by Q_{20}, while $\bar{Q}^*_n, g^*_n$ are the targeted estimators of $\bar{Q}_0$, g_0, using the TMLE.

Given $(\bar{Q}, g)$, the local least favorable submodel through $\bar{Q}$ and local desired submodel through g are defined by

$$\begin{aligned} \text{Logit}\bar{Q}^{lfm}_\epsilon &= \text{Logit}\bar{Q} - \epsilon H_1(g) \\ \text{Logit}\bar{g}^{sm}_\epsilon &= \text{Logit}\bar{g} - \epsilon H_2(\bar{Q}, Q_2, g). \end{aligned}$$

Let $L_2(g) = -\log g$, and $L(\bar{Q})(O) = -\{Y \log \bar{Q} + (1 - Y) \log(1 - \bar{Q})\}$ be the quassi-log-likelihood loss. Let $\bar{L}(\bar{Q}, g) = L(\bar{Q}) + L_2(g)$ be the sum loss function for $(\bar{Q}, g)$. The corresponding universal score-specific submodel through $(\bar{Q}, g)$ is defined by the following differential recursive relation: for $\epsilon > 0$

$$\begin{aligned} \text{Logit}\bar{Q}_{\epsilon+d\epsilon} &= \text{Logit}\bar{Q}_{\epsilon} - d\epsilon H_1(g_{\epsilon}) \\ \text{Logit}\bar{g}_{\epsilon+d\epsilon} &= \text{Logit}\bar{g}_{\epsilon} - d\epsilon H_2(\bar{Q}_{\epsilon}, Q_2, g_{\epsilon}). \end{aligned}$$

Similarly, we can define this submodel for $\epsilon < 0$. Equivalently, their integral representation is given by: for $\epsilon > 0$

$$\begin{aligned} \text{Logit}\bar{Q}_{\epsilon} &= \text{Logit}\bar{Q} - \int_0^{\epsilon} H_1(g_x)dx \\ \text{Logit}\bar{g}_{\epsilon} &= \text{Logit}\bar{g} - \int_0^{\epsilon} H_2(\bar{Q}_x, Q_2, g_x)dx, \end{aligned}$$

and, for $\epsilon < 0$,

$$\begin{aligned} \text{Logit}\bar{Q}_{\epsilon} &= \text{Logit}\bar{Q} + \int_{\epsilon}^{0} H_1(g_x)dx \\ \text{Logit}\bar{g}_{\epsilon} &= \text{Logit}\bar{g} + \int_{\epsilon}^{0} H_2(\bar{Q}_x, Q_2, g_x)dx, \end{aligned}$$

The TME based on this universal score-specific submodel is now computed as follows. Let $Q_{W,n}, \bar{Q}_n, g_n, Q_{2n}$ be the initial estimators. Let $h > 0$ be a small number. Determine first in which direction the empirical risk increases: $P_n\bar{L}(\bar{Q}_{n,h}, g_{n,h}) < P_n\bar{L}(\bar{Q}_n, g_n)$ or $P_n\bar{L}(\bar{Q}_{n,-h}, g_{n,-h}) < P_n\bar{L}(\bar{Q}_n, g_n)$. Suppose that $h > 0$ is the direction that decreases the empirical risk of the sum loss function. Now, one finds the first local minimum ϵ_n of $\epsilon \rightarrow P_n\bar{L}(\bar{Q}_{n,\epsilon}, g_n^* = g_{n,\epsilon})$ for $\epsilon > 0$. The TMLE of $(Q_{W,0}, \bar{Q}_0, g_0, Q_{20})$ using this universal score-specific submodel is defined by $(Q_{W,n}, \bar{Q}_n^* = \bar{Q}_{n,\epsilon_n}, g_{n,\epsilon_n}, Q_{2n})$, and the corresponding TMLE of ψ_0 is given by $\Psi(Q_{W,n}, \bar{Q}_{n,\epsilon_n})$. The TMLE solves $P_n\{D^*(Q_{W,n}, \bar{Q}_n^*, g_n^*) + D_2(\bar{Q}_n^*, Q_{2n}, g_n^*)\} = 0$. By definition of D_2, the correction D_2 improves the efficiency of the TMLE relative to the TMLE that does not use this correction.

C.2 Using a universal score-specific submodel to obtain asymptotic linearity under milder conditions

Consider again the setting that $O \sim P_0 \in \mathcal{M}$, $\Psi(P) = \Psi_1(Q(P))$, $D^*(P) = D^*(Q(P), G(P))$ for a nuisance parameter $G(P)$ orthogonal to $Q(P)$. In the previous subsection we showed that targeting an initial estimator g_n can make the TMLE more efficient at misspecified Q_n when g_n is a well behaved MLE of g_0 under a correctly specified model $\mathcal{G}$ for g_0.

Suppose now that g_n is based on a machine learning algorithm such as the ensemble super-learner based on a user supplied library of machine learning algorithms. We want to guarantee that the TMLE remains asymptotically linear even when Q_n is misspecified, but now without relying on g_n to be an MLE of a relatively small correct model. Instead we will rely on g_n to converge at a good enough (non-$\sqrt{n}$-rate) to g_0 (van der Laan, 2012a). We will now show how this can be achieved with a universal score-specific submodel. Suppose that we use the TMLE $(Q_n^* = Q_{n,\epsilon_n}, G_n^* = G_{n,\epsilon_n})$ based on a universal score-specific submodel $(Q_{n,\epsilon}, G_{n,\epsilon})$ so that

$$P_n\{D^*(Q_n^*, G_n^*) + D_2(Q_n^*, G_n^*)\} = 0. \tag{14}$$

We will now go through a template for proving asymptotic linearity of $\Psi(Q_n^*)$, which will then demonstrate how D_2 needs to be chosen. Firstly, we use that

$$-P_0 D^*(Q_n^*, G_n^*) = \Psi(Q_n^*) - \Psi(Q_0) + R_2(Q_n^*, Q_0, G_n^*, G_0), \tag{15}$$

where R_2 is a second order term in differences $(f_1(Q_n^*) - f_1(Q_0))$ and $f_2(G_n^*) - f_2(G_0)$ for some f_1, f_2. Since Q_n^* can be inconsistent, this second order term cannot be assumed to be negligible. This second order term is assumed to have the so called double robust structure so that $R_2(Q_0, Q_0, G, G_0) = R_2(Q, Q_0, G_0, G_0) = 0$, i..e it equals zero when either Q_0 or G_0 is correctly specified. Combining (14) and (19) yields:

$$\begin{aligned}(P_n - P_0)\{D^*(Q_n^*, G_n^*) + D_2(Q_n^*, G_n^*)\} &= \Psi(Q_n^*) - \Psi(Q_0) \\ &\quad + R_2(Q_n^*, Q_0, G_n^*, G_0) - P_0 D_2(Q_n^*, G_n^*).\end{aligned} \tag{16}$$

Suppose now that by utilizing the special structure of $R_2()$ we can construct a data adaptive real valued $G \rightarrow \Phi_n(G)$ such that

$$R_2(Q_n^*, Q_0, G_n^*, G_0) = \Phi_n(G_n^*) - \Phi_n(G_0) + R_{2n}^a, \tag{17}$$

for some second order term in terms of differences $(G_n - G_0)$ and $Q_n^r - Q_0^r$ for some much easier to estimate parameter Q_0^r of (Q_0, G_0). We would now assume that $R_{2n}^a = o_P(1/\sqrt{n})$.

For example, in the EY_1 example, we have

$$\begin{aligned}R_2(Q, Q_0, G, G_0) &= P_0(\bar{Q} - \bar{Q}_0)(\bar{g} - \bar{g}_0)/\bar{g}) \\ &= E_0 E_0(Y - \bar{Q}(W) \mid A = 1, \bar{g}_0, \bar{Q})(\bar{g} - \bar{g}_0)/\bar{g}) \\ &= \Phi_{\bar{Q},\bar{g},\bar{g}_0,0}(\bar{g}) - \Phi_{\bar{Q},\bar{g}_0,0}(\bar{g}_0),\end{aligned}$$

where

$$\Phi_{\bar{Q},\bar{g},\bar{g}_0,0}(\bar{g}_1) = \int E_0(Y - \bar{Q} \mid A = 1, \bar{g}_0, \bar{Q})\bar{g}_1/\bar{g} dP_0(w).$$

Define $Q_{20} = E_0(Y - \bar{Q} \mid A = 1, \bar{g}_0, \bar{Q})$ and let Q_{2n} be the corresponding estimator $E_n(Y - \bar{Q}_n \mid A = 1, \bar{g}_n, \bar{Q}_n)$, treating $\bar{g}_n, \bar{Q}_n$ as fixed functions of W. Then, we can also denote $\Phi_{\bar{Q},\bar{g},\bar{g}_0,0} = \Phi_{Q_{20},Q_{W,0}}$, and we can define Φ_n by $\Phi_{Q_{2n},Q_{W,n}}$. Thus, in this example, we can define

$$\Phi_n(\bar{g}_1) = E_{P_n} E_n(Y - \bar{Q}_n \mid A = 1, \bar{g}_n, \bar{Q}_n)\bar{g}_1/\bar{g}_n,$$

and the second order term R^a_{2n} involves square differences $(\bar{g}_n-\bar{g}_0)^2$, $(\bar{Q}_{2n}-\bar{Q}_{20})(\bar{g}_n-\bar{g}_0)$, and square differences involving (P_n-P_0) over W, all reasonable second order terms.

So combining (16) with (17) yields:

$$\begin{aligned}(P_n-P_0)\{D^*(Q_n^*,G_n^*)+D_2(Q_n^*,G_n^*)\} &= \Psi(Q_n^*)-\Psi(Q_0)\\ &\quad +\Phi_n(G_n^*)-\Phi_n(G_0)-P_0D_2(Q_n^*,G_n^*)+o_P(1/\sqrt{n}).\end{aligned}$$

Let $D_{2,n}(P_0)$ be efficient influence curve of Φ_n at P_0, viewing Φ_n as a given real value parameter defined on $\mathcal{M}$. By augmenting the original definition of Q with whatever extra parameters are needed to evaluate this efficient influence curve, we can denote $D_{2,0}(P_0)$ with $D_2(Q_0,G_0,\gamma_0)$, where γ_0 is the part that is externally estimated with γ_n. By the general property of an canonical gradient of a target parameter mapping, one will have that

$$-P_0D_2(Q_n^*,G_n^*,\gamma_n)=\Phi_n(G_0)-\Phi_n(G_n)+R^b_{2n}, \tag{18}$$

where R^b_{2n} is a second order term. We will assume $R^b_{2n}=o_P(1/\sqrt{n})$. Combining this with the previous equation yields:

$$(P_n-P_0)\{D^*(Q_n^*,G_n^*)+D_2(Q_n^*,G_n^*,\gamma_n)\}=\Psi(Q_n^*)-\Psi(Q_0)+o_P(1/\sqrt{n}),$$

where the $o_P(1/\sqrt{n})$ now equals $R^a_{2n}+R^b_{2n}$. That is, we have shown

$$\Psi(Q_n^*)-\Psi(Q_0)=(P_n-P_0)\{D^*(Q_n^*,G_n^*)+D_2(Q_n^*,G_n^*,\gamma_n)\}+o_P(1/\sqrt{n}).$$

We can now finalize the proof as usual by assuming that $\bar{D}(Q_n^*,G_n^*)=D^*(Q_n^*,G_n^*)+D_2(Q_n^*,G_n^*,\gamma_n)$ falls in a P_0-Donsker class with probability tending to 1, and $P_0\{\bar{D}(Q_n^*,G_n^*,\gamma_n)-\bar{D}(Q,G_0,\gamma_0)\}^2$ converges to zero in probability for some possibly misspecified $Q\neq Q_0$, so that

$$\Psi(Q_n^*)-\Psi(Q_0)=(P_n-P_0)\bar{D}(Q,G_0,\gamma_0)+o_P(1/\sqrt{n}).$$

When $Q=Q_0$, it follows that $D_2(Q_0,G_0,\gamma_0)=0$, so that this TMLE $\Psi(Q_n^*)$ is asymptotically efficient when both Q_n,G_n are consistent.

To conclude, we selected $D_2(Q_0,G_0,\gamma_0)$ to be equal to the efficient influence curve of $G\rightarrow\Phi_0(G)$, a parameter that is constructed by careful study of the second order term $R_2(Q,Q_0,G,G_0)\approx\Phi_0(G)-\Phi_0(G_0)$ where the dependence on P_0 of Φ_0 requires a much easier to estimate function of Q_0,G_0. Using the TMLE based on the corresponding universal score-specific submodel, we obtain a TMLE that preserves asymptotic linearity when Q_n is inconsistent, but still consistent for the easier to estimate pieces needed to make the second order terms, R^a_{2n},R^b_{2n}, $o_P(1/\sqrt{n})$, under regularity conditions.

The proof above proves the following formal theorem.

Theorem 6 *Define the second order term $R_2()$ by*

$$-P_0 D^*(Q,G) = \Psi(Q) - \Psi(Q_0) + R_2(Q,Q_0,G,G_0). \qquad (19)$$

For a given (Q_1, G_1, γ), let $\Phi_{Q_1,G_1,\gamma} : \mathcal{M} \rightarrow \Phi_{Q_1,G_1,\gamma}(\mathcal{M})$ be a parameter mapping, where $\Phi_{Q_1,G_1,\gamma}(P) = \Phi_{1,Q_1,G_1,\gamma}(G(P))$ only depends on P through $G(P)$, and it is indexed by an unknown parameter $\Gamma : \mathcal{M} \rightarrow \Gamma(\mathcal{M})$ (which can be consistently estimated). We use this parameter to approximate the second order term $R_2()$ as follows:

$$R_2(Q,Q_0,G,G_0) = \Phi_{1,Q,G,\gamma_0}(G) - \Phi_{1,Q,G,\gamma_0}(G_0) + R_2^a(\gamma_0, Q^r, Q_0^r, G, G_0)$$

for some second order term R_2^a in differences $Q^r - Q_0^r$ and $G - G_0$ for some relatively easy to estimate Q_0^r (relative to original Q_0). Let $D_{2,Q,G,\gamma}(Q_0, G_0)$ be the efficient influence curve of $\Phi_{Q,G,\gamma}$ at P_0. Let the second order term $R_{2,Q,G,\gamma}()$ be defined by:

$$-P_0 D_{2,Q,G,\gamma}(Q,G) = \Phi_{Q,G,\gamma}(G_0) - \Phi_{Q,G,\gamma}(G) + R_{2,Q,G,\gamma}(Q^r, Q_0^r, G, G_0),$$

where again $R_{2,Q,G,\gamma}()$ is second order in terms of an easier to estimate parameter Q_0^r instead of original Q_0.

Let γ_n be a consistent estimator of γ_0. Let (Q_n^*, G_n^*) be an estimator of (Q_0, G_0) that solves

$$0 = P_n \bar{D}(Q_n^*, G_n^*, \gamma_n) \equiv P_n\{D^*(Q_n^*, G_n^*) + D_{2,Q_n^*,G_n^*,\gamma_n}(Q_n^*, G_n^*)\}.$$

Assume $R_2^a(\gamma_n, Q_{r,n}^*, Q_0^r, G_n^*, G_0) = o_P(1/\sqrt{n})$ and $R_{2,Q_n^*,G_n^*,\gamma_n}(Q_n^{r*}, Q_0^r, G_n^*, G_0) = o_P(1/\sqrt{n})$. Assume also that $\bar{D}(Q_n^*, G_n^*, \gamma_n)$ falls in a P_0-Donsker class with probability tending to 1, $P_0\{\bar{D}(Q_n^*, G_n^*, \gamma_n) - \bar{D}(Q, G_0, \gamma_0)\}^2$ converges to zero in probability for some possibly misspecified $Q \neq Q_0$. Then,

$$\Psi(Q_n^*) - \Psi(Q_0) = (P_n - P_0)\bar{D}(Q, G_0, \gamma_0) + o_P(1/\sqrt{n}).$$

C.3 Universal score-specific submodels for one-step higher-order TMLE

Of course, the above formulation can be further generalized as follows. Given a local desired submodel for which $\frac{d}{d\epsilon} L(Q_\epsilon^{lfm})\Big|_{\epsilon=0} = D(Q,G)$ for some specified $D(Q,G)$, the corresponding universal score-specific submodel is defined by the recursive differential equation definition:

$$Q_{\epsilon+d\epsilon} = Q_{\epsilon,d\epsilon}^{lfm}.$$

Under weak regularity condition, this now satisfies that $\frac{d}{d\epsilon} L(Q_\epsilon) = D(Q_\epsilon, G)$, and the one-step TMLE defined by Q_{ϵ_n} with $\epsilon_n = \arg\min_\epsilon P_n L(Q_\epsilon)$ solves $P_n D(Q_{\epsilon_n}, G) = 0$. Therefore, this universal score-specific submodel can also be used to define one-step second-order TMLE of second order pathwise differentiable parameters(Carone et al., 2014; Diaz et al., 2015). In this case $D(Q,G)$ plays the role of $D(Q,G) = D^1(Q,G)+P_nD^2(Q,G)$, where $D^j(Q,G)$ is the j-th order efficient influence function, $j=1,2$. Given an initial estimator (Q_n,G_n), the TMLE Q_{n,ϵ_n} solves $P_nD^1(Q_{n,\epsilon_n},G_n)+P_n^2D^2(Q_{n,\epsilon_n},G_n)=0$, providing the basis for asymptotic efficiency of the second order TMLE under a condition that a third-order difference between (Q_{ϵ_n},G_n) and (Q_0,G_0) is $o_P(1/\sqrt{n})$, while a first order TMLE relies on a second order difference being $o_P(1/\sqrt{n})$.

D Generalization to universal least favorable submodels with loss-functions that depend on nuisance parameters

This section could be read after Section D Let $O \sim P_0 \in \mathcal{M}$, $\Psi : \mathcal{M} \rightarrow \mathbb{R}$, $D^*(P) = D^*(Q(P),G(P))$, $G(P)$ is orthogonal to $Q(P)$. Consider a loss function $L_\Gamma(Q)$ so that $Q(P) = \arg\min_Q PL_{\Gamma(P)}(Q)$, where $\Gamma : \mathcal{M} \rightarrow \Gamma(\mathcal{M})$ is some nuisance parameter. For example, $\Gamma(P)$ might depend on P through $Q(P)$, $G(P)$, or both $(Q(P),G(P))$. Let $\{Q_\epsilon^{lfm} : \epsilon\}$ be a local least favorable submodel through $Q=Q(P)$ at $\epsilon = 0$ w.r.t. this loss function L_γ:

$$\frac{d}{d\epsilon}L_{\Gamma(P)}(Q_\epsilon^{lfm})\bigg|_{\epsilon=0} = D^*(Q(P),G(P)).$$

A TMLE based on this *local* least favorable submodel could now proceed in the following two manners. Simultaneously, the resulting universal least favorable submodel and corresponding one-step TMLE will follow naturally and be described as well.

Case I: Fixing the nuisance parameter in the loss-function. Given an initial (Q,G), and corresponding $\gamma = \Gamma(Q,G)$ or external estimate γ, one defines

$$\epsilon_n^0 = \arg\min_\epsilon P_nL_\gamma(Q_\epsilon^{lfm}),$$

one defines the update $Q^1 = Q_{\epsilon_n^0}^{lfm}$, and one iterates this updating process with

$$\epsilon_n^k = \arg\min_\epsilon P_nL_\gamma(Q_\epsilon^{k,lfm}),$$

$k=1,2,\ldots$ till $\epsilon_n^K \approx 0$, thus fixing γ throughout. The TMLE of Q_0 based on this local least favorable submodel is now $Q^* = Q^K$, and solves

$$P_nD^*(\gamma,Q^*,G) \approx 0,$$

where

$$D^*(\gamma,Q,G) = \frac{d}{d\epsilon}L_\gamma(Q_\epsilon^{lfm})\bigg|_{\epsilon=0}.$$

Under reasonable conditions on the estimator of $\gamma_0 = \Gamma(P_0)$, one will still have

$$-P_0 D^*(\gamma, Q, G) = \Psi(Q) - \Psi(Q_0) + R_2^{\#}(\gamma, \gamma_0, Q, Q_0, G, G_0), \quad (20)$$

for a second order term involving square differences of $(Q - Q_0)$, $G - G_0$, and $\gamma - \gamma_0$. Therefore, one can still establish asymptotic efficiency of such a TMLE under the condition that the second order term is $o_P(1/\sqrt{n})$, and some regularity conditions. The price we paid by fixing the nuisance parameter in the loss function is that the TMLE now solves an incompatible efficient influence curve equation in the sense that the estimator γ will not be compatible with the TMLE (Q_n^*, G_n). Generally, speaking this seems of little consequence, as long as $D^*(\gamma, Q, G)$ still has the desired second order expansion (20).

The construction of an L_γ-specific universal least favorable submodel can now proceed analogue to the case that the loss-function was known by replacing $L(Q)$ by $L_\gamma(Q)$, and $D^*(Q, G)$ by $D^*(\gamma, Q, G)$ fixing γ. In other words, we define the L_γ-specific universal least favorable submodel by the differential equation: for $\epsilon > 0$ and $d\epsilon > 0$,

$$Q_{\epsilon + d\epsilon} = Q_{\epsilon, d\epsilon}^{lfm},$$

and, similarly for $\epsilon < 0$ and $d\epsilon < 0$, we define $Q_{\epsilon - d\epsilon} = Q_{\epsilon, d\epsilon}^{lfm}$. By our previous results, we now have that for all $\epsilon > 0$,

$$\frac{d}{d\epsilon} L_\gamma(Q_\epsilon) = D^*(\gamma, Q_\epsilon, G),$$

and similarly for $\epsilon < 0$. The TMLE using this L_γ-specific universal least favorable submodel takes only one step so that the TMLE of Q_0 is given by $Q^* = Q_{\epsilon_n^0}$, solving $P_n D^*(\gamma, Q^*, G) = 0$.

Case II: Updating the nuisance parameter. Given an initial (Q, G), and corresponding $\gamma = \Gamma(Q, G)$, one defines

$$\epsilon_n^0 = \arg\min_\epsilon P_n L_\gamma(Q_\epsilon^{lfm}).$$

One defines the update $Q^1 = Q_{\epsilon_n^0}^{lfm}$, and $\gamma^1 = \Gamma(Q^1, G)$, and one iterates this updating process with

$$\epsilon_n^k = \arg\min_\epsilon P_n L_{\gamma^k}(Q_\epsilon^{k,lfm}),$$

$k = 1, 2, \ldots$ till $\epsilon_n^K \approx 0$, thus updating γ^k throughout. The TMLE of Q_0 based on this local least favorable submodel is now $Q^* = Q^K$, and solves

$$P_n D^*(Q^*, G) \approx 0.$$

The asymptotic efficiency of the TMLE under the usual conditions follows accordingly.

We define the universal least favorable submodel by the same differential equation as above for the fixed loss-function case: for $\epsilon > 0$ and $d\epsilon > 0$,

$$Q_{\epsilon+d\epsilon} = Q^{lfm}_{\epsilon,d\epsilon},$$

and, similarly for $\epsilon < 0$ and $d\epsilon < 0$, we define $Q_{\epsilon-d\epsilon} = Q^{lfm}_{\epsilon,d\epsilon}$. As a consequence, for all $\epsilon > 0$,

$$\frac{d}{dh} L_{\Gamma(Q_\epsilon,G)}(Q_{\epsilon+h})\bigg|_{h=0} = D^*(Q_\epsilon, G).$$

Thus, for all $\epsilon > 0$,

$$\frac{d}{d\epsilon} L_\gamma(Q_\epsilon) = D^*(Q_\epsilon, G).$$

The MLE-step for the one-step TMLE is now defined as follows. First determine the sign of h for which $P_n L_{\gamma(Q,G)}(Q^{lfm}_{dh}) < P_n L_{\gamma(Q,G)}(Q)$. Suppose the empirical risk decreases in the direction $h > 0$. Now, we determine the first $\epsilon^0_n > 0$ for which

$$\frac{d}{dh} P_n L_{\Gamma(Q_\epsilon,G)}(Q_{\epsilon+h})\bigg|_{h=0} = 0,$$

or equivalently, at which

$$P_n D^*(Q_\epsilon, G) = 0.$$

Notice that this corresponds with the first ϵ^0_n at which $P_n L_{\gamma_{\epsilon^0_n}}(Q_{\epsilon^0_n+h})$ is not increasing in $h > 0$ anymore.

The TMLE using this universal least favorable submodel w.r.t loss $L_\gamma(Q)$ takes only one step so that the TMLE of Q_0 is given by $Q^* = Q_{\epsilon^0_n}$, solving $P_n D^*(Q^*, G) = 0$.

D.1 Example: Sequential regression TMLE of counterfactual mean for multiple time point intervention using universal least favorable model

Here we develop a TMLE based on the universal one-dimensional least favorable submodel, while in our previous work (Gruber and van der Laan, 2012; Bang and Robins, 2005) we use a local least favorable submodel with a parameter for each time point. Let $O = (L(0), A(0), L(1), A(1), Y) \sim P_0$, and let the statistical model $\mathcal{M}$ only put restrictions on the conditional probability distributions $g_{A(0)}$ and $g_{A(1)}$ of $A(0)$, given $L(0)$, and $A(1)$, given $\bar{L}(1), A(0)$, respectively. Let $L(0) \to d_0(L(0))$ and $\bar{L}(1) \to d_1(\bar{L}(1))$ be two functions that can be used to deterministically assign treatment $A(0) = d_0(L(0))$ and $A(1) = d_1(\bar{L}(1))$, respectively. Let $\bar{d} = (d_0, d_1)$. Given this dynamic treatment regimen (d_0, d_1) we define the target parameter $\Psi : \mathcal{M} \to \mathbb{R}$ by

$$\Psi(P) = E_P(E_P(E_P(Y \mid \bar{A}(1) = \bar{d}(\bar{L}(1)), \bar{L}(1)) \mid A(0) = d_0(L(0)), L(0))).$$

Let $\bar{Q}^2 = E_P(Y \mid \bar{A}(1) = \bar{d}(\bar{L}(1)), \bar{L}(1))$, $\bar{Q}^1 = E_P(\bar{Q}^2 \mid A(0) = d_0(L(0)), L(0))$, and $\bar{Q}^0 = E_P(\bar{Q}^1)$. Let $\bar{Q} = (\bar{Q}^2, \bar{Q}^1)$, and $Q = (\bar{Q}, \bar{Q}^0)$ and note that $\Psi(P) = \Psi_1(Q) = \bar{Q}^0$.

The efficient influence curve of Ψ at P is given by:

$$\begin{aligned} D^*(P) &= \{\bar{Q}^1 - \bar{Q}^0\} + \frac{I(A(0) = d_0(L(0)))}{g_{A(0)}(O)}(\bar{Q}^2 - \bar{Q}^1) \\ &\quad + \frac{I(\bar{A}(1) = \bar{d}(\bar{L}(1)))}{g_{A(0)} g_{A(1)}(O)}(Y - \bar{Q}^2) \\ &\equiv D_0^*(P) + D_1^*(P) + D_2^*(P) \end{aligned}$$

We will also denote $D^*(P)$ with $D^*(Q, g)$, $g = (g_{A(0)}, g_{A(1)})$, and $D_2^*(P) = D_2^*(\bar{Q}, g)$, $D_1^*(P) = D_1^*(\bar{Q}, g)$.

Consider the following loss functions for the components of $\bar{Q}$:

$$\begin{aligned} L_2(\bar{Q}^2) &= -I(\bar{A}(1) = \bar{d}(\bar{L}(1)))\{Y \log \bar{Q}^2 + (1 - Y)\log(1 - \bar{Q}^2)\} \\ L_{1,\bar{Q}^2}(\bar{Q}^1) &= -I(A(0) = d_0(L(0)))\{\bar{Q}^2 \log \bar{Q}^1 + (1 - \bar{Q}^2)\log(1 - \bar{Q}^1)\}. \end{aligned}$$

Given $\bar{Q}_n^1$, we will estimate $\bar{Q}_0^0$ with $\bar{Q}_n^0 = P_n \bar{Q}_n^1$, an empirical mean. As a consequence, we only need a TMLE of $\bar{Q}_0^2$ and $\bar{Q}_0^1$, and the TMLE of $\bar{Q}_0^0$ follows by taking the empirical mean over $L(0)$ of the TMLE of $\bar{Q}_0^1$.

We can now define the sum loss function for $\bar{Q}$:

$$L_{\bar{Q}^{2,\prime}}(\bar{Q}) \equiv L_2(\bar{Q}^2) + L_{1,\bar{Q}^{2\prime}}(\bar{Q}^1),$$

which is indexed by nuisance parameter $\bar{Q}^2$ itself. For notational convenience, let's denote this nuisance parameter with $\Gamma(\bar{Q}) = \bar{Q}^2$. Then, this loss-function can also be represented as:

$$L_\gamma(\bar{Q}^2, \bar{Q}^1) = L_2(\bar{Q}^2) + L_{1,\gamma}(\bar{Q}^1).$$

Indeed, we have $L_{\gamma_0}(\bar{Q})$ is a valid loss function for $\bar{Q}_0 = \arg\min_{\bar{Q}} P_0 L_{\gamma_0}(\bar{Q})$. Consider the following local least favorable submodel through $\bar{Q}$:

$$\begin{aligned} \text{Logit}\bar{Q}_\epsilon^{2,lfm} &= \text{Logit}\bar{Q}^2 - \epsilon H_2(g) \\ \text{Logit}\bar{Q}_\epsilon^{1,lfm} &= \text{Logit}\bar{Q}^1 - \epsilon H_1(g) \end{aligned}$$

where $H_2(g) = I(\bar{A}(1) = \bar{d}(\bar{L}(1)))/(g_{A(0)} g_{A(1)}(O))$ and $H_1(g) = I(A(0) = d_0(L(0)))/g_{A(0)}(O)$. Indeed, we have

$$\left.\frac{d}{d\epsilon} L_{\bar{Q}^2}(\bar{Q}_\epsilon^{lfm})\right|_{\epsilon=0} = D^{*2}(\bar{Q}, g) + D^{*1}(\bar{Q}, g).$$

Case I: Fixed loss function L_γ. The corresponding L_γ-specific universal least favorable submodels are defined by the differential equation $\bar{Q}^2_{\epsilon+d\epsilon} = \bar{Q}^{2,lfm}_{\epsilon,d\epsilon}$ and $\bar{Q}^1_{\epsilon+d\epsilon} = \bar{Q}^{1,lfm}_{\epsilon,d\epsilon}$, which implies the integral representation given by

$$\begin{aligned} \text{Logit}\bar{Q}^2_\epsilon &= \text{Logit}\bar{Q}^2 - \epsilon H_2(g) \\ \text{Logit}\bar{Q}^1_\epsilon &= \text{Logit}\bar{Q}^1 - \epsilon H_1(g). \end{aligned}$$

Thus, the L_γ-specific universal least favorable submodel through $\bar{Q} = (\bar{Q}^2, \bar{Q}^1)$ equals the local least favorable submodel: $\bar{Q}^{lfm}_\epsilon = \bar{Q}_\epsilon$. Indeed,

$$\begin{aligned} \frac{d}{d\epsilon} L_{\bar{Q}^2}(\bar{Q}_\epsilon) &= H_2(g)(Y - \bar{Q}^2_\epsilon) + H_1(g)(\bar{Q}^2 - \bar{Q}^1_\epsilon) \\ &\equiv D^*_2(\bar{Q}^2_\epsilon, g) + D^*_1(\bar{Q}^2, \bar{Q}^1_\epsilon, g). \end{aligned}$$

The one-step TMLE based on this L_γ-specific universal least favorable submodel is defined by

$$\epsilon_n = \arg\min_\epsilon P_n L_{\bar{Q}^2_n}(\bar{Q}_{n,\epsilon}),$$

and the TMLE of $\bar{Q}_0$ is given by $\bar{Q}_{n,\epsilon_n}$. The resulting TMLE of ψ_0 is simply $\Psi(Q_{n,\epsilon_n}) = P_n \bar{Q}^1_{\epsilon_n}$. This TMLE will now solve the incompatible efficient influence curve equation $0 = P_n D^*(\bar{Q}^2_n, Q^*_n, g_n)$ defined by

$$D^*(\bar{Q}^2_n, Q^*_n, g_n) = D^*_2(\bar{Q}^{2*}_n, g_n) + D^*_1(\bar{Q}^2_n, \bar{Q}^{1*}_n, g_n) + D^*_0(\bar{Q}^{1*}_n, \bar{Q}^{0*}_n).$$

The typical TMLE solves the compatible efficient influence curve equation $0 = P_n D^*(Q^*_n, g_n)$, where

$$D^*(Q^*_n, g_n) = D^*_2(\bar{Q}^{2*}_n, g_n) + D^*_1(\bar{Q}^{2*}_n, \bar{Q}^{1*}_n, g_n) + D^*_0(\bar{Q}^{1*}_n, \bar{Q}^{0*}_n).$$

Let's now prove that this incompatible efficient influence curve still allows the desired second order expansion the asymptotic linearity and efficiency proof relies upon. By the general representation theorem for the efficient influence curve in CAR-censored data models (Robins and Rotnitzky, 1992; **?**), we have

$$D^*(Q^*, g) = D_{IPTW}(g, \bar{Q}^{0*}) + D_{CAR}(\bar{Q}^*, g),$$

where $D_{IPTW}(g, \bar{Q}^0) = I(\bar{A} = \bar{d}(\bar{L})/\bar{g}_1 Y - \bar{Q}^0$, and $D_{CAR}(\bar{Q}, g)$ is a score of the censoring mechanism, thereby, being a function of O that has conditional mean zero w.r.t. g (for every value of $\bar{Q}$). Thus the incompatible efficient influence curve $D^*(\bar{Q}^2, Q^*, g)$ can be represented as $D_{IPTW}(g, \bar{Q}^{0*}) + D_{CAR}(\bar{Q}, g)$, where $\bar{Q} \neq \bar{Q}^*$. We have

$$\begin{aligned} P_0 D^*(\bar{Q}^2, Q^*, g) &= P_0\{D_{IPTW}(g, \bar{Q}^{0*}) + D_{CAR}(\bar{Q}, g)\} \\ &= P_0\{D_{IPTW}(g, \bar{Q}^{0*}) + D_{CAR}(\bar{Q}^*, g)\} \\ &\quad + P_0\{D_{CAR}(\bar{Q}^*, g) - D_{CAR}(\bar{Q}, g)\} \\ &= \Psi(Q_0) - \Psi(Q^*) + R_2(Q^*, Q_0, G, G_0) \\ &\quad + P_0\{D_{CAR}(\bar{Q}^*, g) - D_{CAR}(\bar{Q}, g)\}. \end{aligned}$$

So we need to show that the last term is a second order term. But this last term equals:

$$R_{2a}(\bar{Q}^*, \bar{Q}, g, g_0) \equiv P_0\{D_{CAR}(\bar{Q}^*, g) - D_{CAR}(\bar{Q}, g) - D_{CAR}(\bar{Q}^*, g_0) + D_{CAR}(\bar{Q}, g_0)\}.$$

Thus, we conclude that

$$P_0 D^*(\bar{Q}^2, Q^*, g) = \Psi(Q_0) - \Psi(Q^*) + R_2(Q^*, Q_0, g, g_0) + R_{2a}(\bar{Q}^2, \bar{Q}^{2*}, g, g_0),$$

which thus yields a desired double robust second order remainder term defined as the sum of R_2 and R_{2a}. Since the compatible TMLE generates a second order term R_2, it might be the case that for finite samples the second order term $R_2 + R_{2a}$ of the incompatible TMLE is larger.

Case II: Updating the loss function with ϵ. The universal least favorable submodels are defined as above:

$$\begin{aligned} \text{Logit}\bar{Q}^2_\epsilon &= \text{Logit}\bar{Q}^2 - \epsilon H_2(g) \\ \text{Logit}\bar{Q}^1_\epsilon &= \text{Logit}\bar{Q}^1 - \epsilon H_1(g) \end{aligned}$$

Indeed, it has the following key property with respect to the loss function $L_{\bar{Q}^2}(\bar{Q})$:

$$\begin{aligned} \left.\frac{d}{dh} L_{\bar{Q}^2_\epsilon}(\bar{Q}_{\epsilon+dh})\right|_{h=0} &= H_2(g)(Y - \bar{Q}^2_\epsilon) + H_1(g)(\bar{Q}^2_\epsilon - \bar{Q}^1_\epsilon) \\ &\equiv D^*_2(\bar{Q}_\epsilon, g) + D^*_1(\bar{Q}_\epsilon, g). \end{aligned}$$

Let's assume that we determined that the empirical risk $P_n L_{\bar{Q}^2_n}(\bar{Q}_{n,\epsilon})$ is decreasing at $\epsilon = 0$, so that we need to determine the desired $\epsilon_n > 0$. The solution ϵ_n is defined by the smallest $\epsilon > 0$ for which

$$\left.\frac{d}{dh} P_n L_{\bar{Q}^2_{n,\epsilon}}(\bar{Q}_{n,\epsilon+h})\right|_{h=0},$$

or, equivalently, the smallest $\epsilon > 0$ for which

$$P_n D^*(Q_{n,\epsilon}, g_n) = 0,$$

where

$$Q_{n,\epsilon} = (\bar{Q}^2_{n,\epsilon}, \bar{Q}^1_{n,\epsilon}, \bar{Q}^0_{n,\epsilon} = P_n \bar{Q}^1_\epsilon).$$

The TMLE of ψ_0 is now defined by $\Psi(Q_{n,\epsilon_n}) = P_n \bar{Q}^1_{\epsilon_n}$, and it solves $P_n D^*(Q_{n,\epsilon_n}, g_n) = 0$.

To obtain some insight in solving for ϵ_n, note that it requires solving:

$$\begin{aligned} 0 &= \frac{1}{n}\sum_{i=1}^{n}\{\bar{Q}^1_{n,\epsilon}(L_i(0)) - P_n\bar{Q}^1_{n,\epsilon}\} \\ &+ \frac{1}{n}\sum_{i=1}^{n}\frac{I(A_i(0)=d_0(L_i(0)))}{g_{A(0),n}(O_i)}(\bar{Q}^2_{n,\epsilon} - \bar{Q}^1_{n,\epsilon}) \\ &+ \frac{1}{n}\sum_{i=1}^{n}\frac{I(\bar{A}_i(1)=\bar{d}(\bar{L}_i(1)))}{g_{A(0),n}g_{A(1),n}(O_i)}(Y_i - \bar{Q}^2_{n,\epsilon}). \end{aligned}$$

Since $\bar{Q}^j_{n,\epsilon}$ is a simple adjustment of the initial estimator $\bar{Q}^j_n$ (just adding $\epsilon H_j(g_n)$ on the logistic scale), $j = 2, 1$, this estimator is very easy to compute.

This implementation of TMLE is quite different from the current implementation of TMLE that carries out the TMLE update step by fitting a separate ϵ for updating each $\bar{Q}^j$, and sequentially carrying out these updates starting with $\bar{Q}^2$ and going backwards. In addition, it involves first targeting the regression before defining it as outcome for the next regression backwards in time. For example, if there are many treatment nodes over time, then the TMLE presented above still only relies on fitting a single ϵ, while the current TMLE would require iteratively fitting many ϵ_j's. We suspect that the TMLE proposed here could be significantly more stable in finite samples.