TMLE: AS I (DIMLY) UNDERSTAND IT

ROGER KOENKER

1. Introduction

Targeted maximum likelihood estimation is a general class of procedures for estimating semiparametric statistical models, that is models with a nonparametric nuisance parameter. Judea Pearl, in his forward to the comprehensive exposition of van der Laan and Rose (2011), invokes Jacob Marshak's (1952) observations that good policy decisions rarely require a fully specified stochastic model. Instead, decisions can be based on only a few restrictive features of the model, e.g. average treatment effects, estimable without a full stochastic specification.

In these notes I would like to illustrate the application of this principle beginning with some simple missing data and binary treatment settings and concluding with a discussion of a recent paper of Díaz (2015) on quantile treatment effects. This expositional strategy is what might be called "inductive learning:" we start with very simple cases and try to leverage patterns we see into a general strategy. Those who prefer a more hard nosed approach should proceed directly to Appendix A of van der Laan and Rose (2011).

2. Prelude: Missing Data and Double Robustness

It seems appropriate to begin with a brief discussion of Bang and Robins (2005), which can be viewed as a precursor of the TMLE approach. Consider the standard missing data problem. We observe $Z = (Y, M, X)$. The response Y is observed when $M = 1$ and is missing otherwise. The vector of covariates, X, is available and satisfy:

i. $P(M = 1|Y, X) = P(M = 1|X) \equiv \pi(X)$
ii. $\pi(X) > 0 \quad a.e.P.$

The former condition is usually referred to as "missing at random," but perhaps more descriptively as "no unmeasured confounders." The latter condition is usually referred to as the "common support" condition.

Given a random sample of the triples $\{Z_i\}$ $i = 1, ..., n$ we may adopt a parametric model for π, say the logistic, and estimate the unconditional mean of Y using the Horvitz and Thompson (1952) estimator,

$$\hat{\mu} = n^{-1} \sum_{i=1}^{n} M_i Y_i / \hat{\pi}(X_i).$$

In the original survey sampling context, $\pi(X)$ denoted a priori sampling weights and didn't need to be estimated. In the missing data setting $\pi(X)$ is usually referred to as a propensity score and we simply reweight the observed Y_i's by the inverse of their probability of inclusion to the sample.

Alternatively, we may represent μ as,

$$\mu = \mathbb{E}(\mathbb{E}(Y|M = 1, X))$$

and this suggests estimating a model of the form,

$$\mathbb{E}(Y|M = 1, X) = \Psi(s(X, \beta))$$

Version: March 5, 2016. I would like to express my appreciation to Iván Díaz for providing R code for the simulations reported in Díaz and Rosenblum (2015) and Díaz (2015) and which I have adapted somewhat for the present notes, which were prepared for my Spring 2016 topics course on semiparametric methods in econometrics. Read at your own risk.

for some known link function, Ψ and then estimating,

$$\hat{\mu} = n^{-1} \sum_{i=1}^{n} \Psi(s(X_i, \hat{\beta})).$$

Thus, for example, if Ψ is the identity link and $s(X, \beta) = X^\top \beta$, we would simply estimate β on the observed sample and then estimate μ as the mean of the $\hat{Y}_i$'s for the full sample. If Y_i is itself binary we might again use the logistic link, i.e. $\log(u/(1-u))$ and $\hat{\beta}$ would solve the estimating equation,

$$\sum_{i=1}^{n} M_i \nabla_\beta s(X_i, \beta)(Y_i - \Psi(s(X_i, \beta))) = 0,$$

so again $\hat{\beta}$ would be the restricted MLE based on only the $M_i = 1$ observations. When $s(X, \beta) = X^\top \beta$ then $\nabla_\beta s(X_i, \beta) = X$.

Inquiring Minds Would Like to Know

(1) What happened to the reweighting that is usually applied in the logistic estimating equation (IRLS) approach, see e.g. McCullagh and Nelder (1989) or my 508 notes.
(2) Why logit? Why not Cauchit? See e.g. Koenker and Yoon (2009) or my 508 notes.

Rather than choose between these two options, it seems preferable to combine them in some way. How? Ideally, we would like an estimator that consistently estimates our parameter of interest, μ, provided that *either* the propensity score model is correct, or the outcome regression model is correct. Bang and Robins (2005) refer to this property as "double robustness," I would have called it the belt and suspenders principle.

The implementation of this principle is quite straightforward, we simply augment the outcome regression by writing,

$$\mathbb{E}(Y|M = 1, X) = \Psi(s(X, \beta)) + \phi \pi^{-1}(X, \alpha) \equiv \Phi(X, \beta, \phi)$$

Estimation is carried out in two stages: first, α is estimated by logistic regression, then (β, ϕ) are estimated with $\pi^{-1}(X, \hat{\alpha})$ acting as an additional covariate. The latter step solves the estimating equation,

$$0 = \sum_{i=1}^{n} M_i \nabla \Phi(Y_i - \Phi(X_i, \beta, \phi))$$

which can be rewritten as,

$$0 = U(\mu) = \sum_{i=1}^{n} \pi_i^{-1} M_i \nabla \Phi(Y_i - \mu) - (\pi_i^{-1} M_i - 1)(\Phi(X_i, \beta, \phi) - \mu)$$

or

$$0 = U(\mu) = \sum_{i=1}^{n} \pi_i^{-1} M_i \nabla \Phi(Y_i - \Phi_i) - (\Phi_i - \mu).$$

Note that the first term is zero since it is just the ϕ component of the original estimating equation, so,

$$\hat{\mu} = n^{-1} \sum_{i=1}^{n} \Phi(X_i, \hat{\beta}, \hat{\phi}),$$

can be used as a plugin estimator.

The foregoing missing data formulation can be extended to the standard average treatment effect problem. The data structure is almost the same, we have $Z = (Y, T, X)$, with the assumption that $Y \perp\!\!\!\perp T|X$. Now $T = 1$ indicates treated observations and $T = 0$, controls. We are interested in estimating the average treatment effect, which can be represented either as

$$\mu = \mathbb{E}(TY/\pi(X)) - \mathbb{E}((1-T)Y/(1-\pi(X))),$$

or

$$\mu = \mathbb{E}(\mathbb{E}(Y|T = 1, X)) - \mathbb{E}(\mathbb{E}(Y|T = 0, X)).$$

Thus, we can estimate μ either by

$$\hat{\mu} = \sum_{i=1}^{n} (T_i Y_i / \hat{\pi}_i) - ((1 - T_i) Y_i / (1 - \hat{\pi}_i)),$$

or

$$\hat{\mu} = \sum_{i=1}^{n} [\Psi(s(1, X_i, \hat{\beta})) - \Psi(s(0, X_i, \hat{\beta}))],$$

where $\hat{\pi}_i = \pi(X_i, \hat{\alpha})$ is the estimated propensity score for receiving treatment. To combine these approaches we define the augmented outcome model,

$$\Phi(T, X, \beta, \phi) = \Psi(s(T, X, \beta)) + \phi_1 T \pi^{-1} + \phi_0 (1 - T)(1 - \pi)^{-1}.$$

We then estimate (β, ϕ) as before with $\hat{\alpha}$ fixed and use,

$$\hat{\mu} = \sum_{i=1}^{n} [\Phi(s(1, X_i, \hat{\beta}, \hat{\phi})) - \Phi(s(0, X_i, \hat{\beta}, \hat{\phi}))].$$

3. TMLE: A Missing Manual

Thus far we have not mentioned efficiency. Double robustness was originally focused on achieving consistency under weaker conditions than those employed in earlier propensity score literature. For a more extensive treatment of this approach see Kang and Schafer (2007) and the discussion thereof. Rotnitzky and Vansteelandt (2014) and van der Laan and Rubin (2006) address the efficiency objective relying to some degree on the prior work of Robins, van der Vaart and others. Their strategy is to devise one-step and multi-step iterative procedures based on the efficient semiparametric score function. I will try to briefly describe how this is supposed to work for the standard model for missing data, and for the conventional binary treatment effect model. While the exposition in this section focuses entirely on mean effects, have faith we are (always) headed toward quantiles.

As in the previous section, we observe a random sample, $Z_1, \ldots, Z_n$ with $Z = (Y, M, X)$ satisfying the conditions appearing there. The likelihood of a Z can be factored as,

$$\mathcal{L}(Z) = P_Y(Y|M, X) P_M(M|X) P_X(X).$$

We wish to make no modeling assumption about the marginal distribution of X, so we will assume throughout that P_X will be replaced by its MLE, the empirical distribution function. In contrast, following convention, we will adopt some convenient parametric binary response model for the propensity score, and denote

$$\pi(X) = P_M(M|X).$$

What distinguishes the TMLE approach is the treatment of P_Y. An initial estimate can be produced based on the $M = 1$ subsample; this may involve some fancy machine learning methods, but I won't delve into that. The crucial aspect of the initial estimate, however it is produced, is that it is inherently suspect. It needs to be modified in some way to account for the fact that it is based entirely on the observed Y data. We have already seen a couple of variants of this idea, but they seem rather ad hoc. It would be nice to have a more principled way to design a modification. Ideally, we would like our estimator, $\hat{\theta}$ of the population mean to be efficient in the sense that it satisfied the estimating equation representing the efficient influence function. Of course, this begs the question of how the EIF can be found, but in the following examples the EIFs will appear like rabbits out of the proverbial hat.

For estimating the population mean in the missing at random model the EIF is given by Bang and Robins (2005) as,

$$D(Z) = \frac{M}{\pi(X)} (Y - \mu(X)) + \mu(X) - \theta(Z).$$

Given this EIF we can consider various one parameter submodels, along the lines of van der Vaart (2000), two are suggested in his Example 25.16:

$$p_t(z) = c(t) \exp(t D(z)) p_0(z),$$

and

$$p_t(z) = (1 + tD(z))p_0(z).$$

A third option is proposed by Díaz and Rosenblum (2015),

$$p_t(z) = c(t)(1 + \exp(-2tD(z)))^{-1}p_0(z),$$

The latter authors propose an iterative scheme for updating our provisional nonparametric model p_0, at each step estimating the auxiliary parameter t, and stopping when the MLE of t occurs sufficiently close to zero. The first option would seem to be the most obvious exponential family choice, the others are intended to deal with the (familiar) difficulty of the unbounded IF. A useful, but rather trivial, exercise would be to verify that the MLE for t yields the EIF when t is evaluated at zero.

I'll illustrate the implementation of this approach in R using the first class of submodels. I presume that it will be obvious how to adapt this code to the other submodels. To get started we need to have data; here is a function to generate data according to the three schemes used in Díaz and Rosenblum (2015) for their simulation experiments. The response Y is binary, and observed only when $M = 1$

```
DGP <- function(n, pMf = 1){
    pMf <- switch(pMf,
        function(X) plogis(1 + 2*X),
        function(X) plogis(-1 + 2*X),
        function(X) plogis(-6 + 2*X + X^2)
        )
    pYf <- function(X) plogis(X - X^2)
    Z <- rnorm(n, sd = 1/sqrt(2))
    X <- Z + rnorm(n, sd = 1/sqrt(2))
    M <- rbinom(n, 1, pMf(X))
    Y <- rbinom(n, 1, pYf(X))
    Y[M==0] <- NA
    data.frame(Y, M, X)
}
```

Next we need code for the log likelihood function and its gradient.

```
logLik <- function(t, py, pm, px, Y, M){
    Dp <- function(py, pm, px)
        function(Y, M) (M/pm) * (Y - py) + py - sum(px * py)
    D <- Dp(py, pm, px)
    A <- cbind(py * pm, (1 - py) * pm, (1 - pm))
    B <- cbind(D(1,1), D(0,1), D(2,0))
    a <- (A * exp(t * B)) %*% rep(1,3)
    -mean(t * D(Y, M)) + log(mean(px * a))
}
gradLik <- function(t, py, pm, px, Y, M){
    Dp <- function(py, pm, px)
        function(Y, M) (M/pm) * (Y - py) + py - sum(px * py)
    D <- Dp(py, pm, px)
    A <- cbind(py * pm, (1 - py) * pm, (1 - pm))
    B <- cbind(D(1,1), D(0,1), D(2,0))
    a <- (A * exp(t * B)) %*% rep(1,3)
    b <- (A * exp(t * B) * B) %*% rep(1,3)
    -mean(D(Y, M)) + mean(px * b)/mean(px * a)
}
```

Finally we have the main TMLE function. The only tricky bit is the way that the function pz gets updated. This is an example of the flexibility of R to "compute on the language," that is for functions to produce functions.

```
tmlex <- function(Z, tol = 1e-4, maxit = 10){
    # Exponential Family TMLE for (binary) missing data model
    # Thanks to Ivan Diaz for original implementation R code
    # Thanks to Mikko Korpela for R-help with the Dcall bit
    pZ <- function(pX, pM, pY)
        function(y,m) pX * pM^m * (1-pM)^(1-m) * pY^(m*y) *
            (1 - pY)^(m * (1 - y)) * (y == 2)^(1-m)
    Dp <- function(py, pm, px)
        function(y, m) (m/pm) * (y - py) + py - sum(px * py)
    Deps <- function(eps, py, pm, px)
        function(y, m) exp(eps * ((m/pm) * (y - py) + py - sum(px * py)))
    Y <- Z$Y; M <- Z$M; X <- Z$X
    Y[is.na(Y)] <- 2  # Aaargh.
    # Initialize
    fitM <- glm(M ~ X + I(X^2), family = binomial, data = Z)
    fitY <- glm(Y ~ X + I(X^2), family = binomial, subset = (M==1), data = Z)
    EM0 <- predict(fitM, newdata = Z, type = "response") # g1X in DR
    EY0 <- predict(fitY, newdata = Z, type = "response") # Q1X in DR
    pz <- pZ(1/nrow(Z), EM0, EY0)
    it <- 0
    eps <- Inf
    while(eps > tol & it < maxit){
        it <- it + 1
        p11 <- pz(1,1); p01 <- pz(0,1); p20 <- pz(2,0)
        px <- p11 + p01 + p20 # P(X)
        pm <- (p11 + p01)/px  # P(M = 1| X)
        py <- p11/(p11 + p01) # P(Y = 1| M = 1, X)
        eps <- optim(0, fn = logLik, gr = gradLik, method = "BFGS",
            py = py, pm = pm, px = px, Y = Y, M = M)$par
        Dcall <- do.call("Deps", list(eps = eps, py = py, pm = pm, px = px))
        body(pz) <- call("*", body(pz), as.call(list(Dcall, quote(y), quote(m))))
    }
    p11 <- pz(1,1); p01 <- pz(0,1); p20 <- pz(2,0)
    px <- p11 + p01 + p20 # P(X)
    pm <- (p11 + p01)/px  # P(M = 1| X)
    py <- p11/(p11 + p01) # P(Y = 1| M = 1, X)
    EIF <- Dp(py, pm, px)
    theta <- sum(px * py)
    list(theta = theta, var = var(EIF(Y,M))/nrow(Z), nit = it)
}
```

Now to illustrate useage we can do:

```
set.seed(1492)
Z <- DGP(1000)
a <- tmlex(Z)
```

Of course the standard MAR model is not that exciting, so it would be nice to demonstrate that there are better motivations for all of this within the domain of econometrics. Fortunately, such a motivation is close at hand, as we have already seen: estimating average treatment effects is essentially the same problem and we can proceed exactly as described at the end of Section 2.

Rather than dwelling on this again, however, I'd like to proceed immediately to something even more compelling – estimation of of quantiles with missing data and quantile treatment effects.

4. Quantiles for Missing Data and QTEs

In this section I'll rely heavily on recent work of Díaz (2015), but let's begin with a brief review of EIFs for quantile models. In the simplest case we have a random sample $\{Y_1, \ldots, Y_n\}$ from a distribution F, and we will assume as usual that F has a density f with respect to Lebesgue measure, and that f is strictly positive on the support of F. Then the ordinary sample quantile,

$$\hat{Q}(\tau) = \inf\{y : \ F_n(y) \geq \tau\}$$

where F_n denotes the usual empirical distribution function $F_n(y) = n^{-1} \sum I(Y_i \leq y)$, satisfies the condition,

$$0 \approx \sum_{i=1}^{n} \psi_\tau(Y_i - \hat{Q}(\tau)),$$

with $\psi_\tau(u) = I(u < 0) - \tau$. The fishy $\approx$ symbol reflects the possible ambiguity when $n\tau$ is an integer and there are multiple solutions, but this isn't very consequential since it is a $1/n$ phenomenon in a $1/\sqrt{n}$ world. A little more effort, see e.g. my 574 Lecture 15, reveals that the efficient influence function for estimating $Q(\tau)$ is given by,

$$\text{EIF}_{\hat{Q}(\tau)}(y) = \frac{1}{f(Q(\tau))} \sum_{i-1}^{n} (I(y < \hat{Q}(\tau)) - \tau).$$

This isn't terribly interesting since the density factor doesn't really play an active role. Things get more interesting in regression settings. We can consider the estimation of the linear quantile regression model,

$$Q_{Y|X}(\tau|x) = x^\top \beta(\tau)$$

as semiparametric with finite dimensional parameter, $\beta(\tau)$ and nuisance parameter, $F_{Y|X}$. Now the EIF for β for fixed τ is

$$\text{EIF}_{\hat{\beta}(\tau)}(y, x) = \sum_{i-1}^{n} \tfrac{1}{f_i(x_i^\top \beta(\tau)))} (I(y < \hat{x}_i^\top \beta(\tau)) - \tau).$$

and each observation is weighted inside the EIF sum according to the local, conditional density terms f_i. This version of the EIF corresponds to the optimally weighted quantile regression problem,

$$\hat{\beta}(\tau) = \text{argmin}_b \sum_{i=1}^{n} f_i(x_i^\top \beta(\tau)) \rho_\tau(y_i - x_i^\top b).$$

The unweighted quantile regression estimator achieves an asymptotic covariance matrix of the sandwich form, $V = H^{-1} J H^{-1}$ where $H = \lim X^\top \Phi X/n$, $\Phi = \text{diag}(f_i)$ and $J = \lim \tau(1-\tau) X^\top X/n$. In the weighted form, the sandwich collapses to $V = \lim \tau(1-\tau) X^\top X/n$. This conforms to applying the f_i weights in the efficient estimating equation and thereby canceling them out in the EIF. A variety of proposals have been made to construct one-step estimators that estimate weights based on a preliminary estimator of $\beta(\tau)$ and reestimate. A possibly interesting alternative to this approach would be to apply the TMLE strategy and iterate toward the solution of the EIF estimating equation. We will leave this to future consideration and now return to the missing data setting, albeit with the objective of estimating quantiles rather than means. Before doing so, I would like to stress that a virtue of the quantile regression paradigm is that even this efficient score formulation depends only on local information about $F_{Y|X}$ near the τth conditional quantile.

The missing data setting for quantile estimation is quite similar to that of the previous sections except that the response, Y, when it is observed is no longer binary. Quantiles for binary data are somewhat silly, although conditional quantiles for binary data have their own charms. We maintain our prior conditions

i. $P(M = 1|Y, X) = P(M = 1|X) \equiv \pi(X)$

ii. $\pi(X) > 0 \quad a.e.P.$
iii. $Y \perp\!\!\!\perp M|X$

Again, we can factor $P_Z = P_Y(y|M,X)P_M(M|X)P_X$ and as before we will denote the propensity score, $\pi(x) = P_M(M = 1|X = x)$. The unconditional distribution of Y will be denoted by F and we will assume that it has strictly positive density in a neighborhood of $\theta = F^{-1}(\tau)$, which will be our target parameter. We will begin by focusing on the pure missing data model, and then turn our attention to the QTE problem. Throughout we will consider the unconditional quantile θ with the understanding that extensions to conditional quantile estimation would be an attractive subject for future research.

The EIF for estimators of θ at P is given by,

$$D(P,Z) = \frac{1}{f(\theta)}\left[\frac{M}{\pi(X)}\left(I(Y < \theta) - P_Y(\theta|1,X)\right) + P_Y(\theta|1,X) - \tau\right].$$

This result, Lemma 1 of Díaz (2015), combines what we have already seen regarding the MAR model for means with the EIF result for unconditional quantiles.

Unlike the binary response setting, or for that matter estimating averages for continuous response, for quantiles we require an initial estimator of the entire conditional distribution, $P_Y(y|M = 1,X)$. In principle this could be accomplished by some form of nonparametric density estimation, but when the dimension of X is large, or even moderate, this is not very practical; so we resort to – what else? – quantile regression. On a grid of $\tau \in (0,1)$, we compute,

$$\hat{Q}_{Y|M=1,X}(\tau|x) = x^\top \hat{\beta}(\tau),$$

and to ensure monotonicity at all x_i of $\hat{Q}$ we "rearrange" $\hat{Q}(\tau|x)$ as in Chernozhukov et al. (2009). The gory details of this process are illustrated in the R code that is provided below for the QTET problem. Our initial estimate of the conditional distribution of Y is easily obtained by inversion, and the marginal by integrating out over the empirical distribution of X. Note that the latter step integrates over the entire sample, not just the subsample with observed response, that is,

$$\hat{F}_n(y) = n^{-1}\sum_{i=1}^{n} \hat{Q}^{-1}(y|M = 1,X_i).$$

Our initial estimator of θ becomes $\hat{\theta} = \hat{F}_n^{-1}(\tau)$.

To update $P_Y(y|M = 1,X)$ Díaz (2015) suggests the same exponential family submodel that we have described above. Let $p_Y(y|M = 1,x)$ be the density corresponding to $P_Y(y|M = 1,x)$, and denote,

$$p_t(y|M = 1,x) = c(t,p_Y)\exp(tD(p_Y,z))p_Y(y|M = 1,x)$$

where $c(t,p_Y)$ is the normalizing constant, and

$$D(p_Y,z) = \frac{1}{\pi(x)}\left[I(y \le \theta) - P_Y(y|M = 1,x)\right].$$

At each iteration we can solve for the MLE of t, and then update to get $p_{\hat{t}}$, continuing until $\hat{t} \approx 0$, at which point we have a fixed point solution for the least favorable model. Rather than dwell further on the missing data case I'd rather proceed, as does Díaz (2015) to consider the closely related case of estimating QTEs.

In the familiar Neyman-Rubin-Pearl causal model we have a binary treatment, T and observe a response Y_1 if $T = 1$, and a response Y_0 if $T = 0$, but never both. We would like to estimate the quantile treatment effect on the treated, by which we will mean the difference between the τth quantile of the response distribution of the treated subjects minus the τth quantile of what would have been the response distribution for the treated subjects if they had not received the treatment. The former quantity is obviously estimable as,

$$\hat{\theta}_1 = \operatorname{argmin}\sum_{i:T_i=1} \rho_\tau(Y_i - \theta),$$

so we will concentrate on the more challenging task of estimating θ_0, the counter-factual τth quantile for the treatment subjects had they been controls. As usual we need to make identifying

assumptions. On the face of it, the task seems formidable since we don't observe *any* Y_0's for the treatment subjects. But covariates come to the rescue. The covariates are the pixie dust, and

i. $P(T = 1|Y, X) = P(T = 1|X) \equiv \pi(X)$
ii. $\pi(X) > 0 \quad a.e.P.$

are the "happy thoughts" that keep you airborne in the night sky of causal inference.

The distribution of the control responses for the treatment group can then be identified from the expression,

$$F^0(y) = \sum_x P_Y(y|T = 0, x) p_X(x|T = 1).$$

Clearly, $P_Y(y|T = 0, x)$ is estimable from the control observations in exactly the same way we were able to estimate $P_Y(y|T = 1, x)$ from the treatment group. So to obtain an estimate of F^0 we simply reweight according the mass associated with the covariates in the treatment group. For example, if there are a lot of $x = a$ types in the treatment group relative to their abundance in the control group, then those guys will get more weight in this sum. This brings us back to something similar to the missing data problem we have already considered.

The EIF for estimating θ_0 is given by,

$$D_0(p_Y, z) = \frac{1}{f(\theta_0}\left[\frac{1-T}{\mathbb{E}T}\frac{\pi(x)}{1-\pi(x)}\left(I(y \le \theta_0) - P_Y(\theta_0|M = 0, x)\right) + \frac{T}{\mathbb{E}T}\left(P_Y(\theta_0|M = 0, x) - \tau\right)\right].$$

This looks rather complicated, but it also closely resembles the pure missing data problem we have already considered, except that we are now focused entirely on the missing control observations for the treatment group.

For those, like me, who believe in mathematics only insofar as it can be translated into R code to compute something, and be experimented with, I offer my slightly modified version of the Díaz (2015) TMLE code for the QTET estimator. It will be noted that the function `QTET` actually returns several competing estimators of the QTET. In addition to the TMLE there is a naive estimator based on only the initial estimators of the quantiles, and two one-step methods, one attributable to Firpo (2007) and the other discussed briefly by Díaz (2015). I will come back to these after a brief discussion of the TMLE implementation.

```
QTET <- function(Z, tau = 0.5, taus = 1:99/100, maxit = 10) {
    # Computes QTET estimators as in Diaz (2015)
    #
    # Args:
    # Z ~ data.frame consisting of:
    # y ~ a response vector
    # t ~ a treatment indicator
    # x ~ some other conditioning covariates
    # tau ~ quantile of interest
    # taus ~ grid of quantiles for CQF estimation
    # maxit ~ upper bound on iteration count
    #
    # Returns:
    # tmle ~ the TMLE estimator of Diaz (2015)
    # se ~ standard error of the tmle
    # firpo ~ Firpo (2007) propensity score estimator
    # onestep ~ one-step estimator based on tmle EIF
    # Some required functions:
    F <- function(x, Q, w) mean(rowSums((Q <= x) * w))
    Finv <- function(u, Q, w, range)
        uniroot(function(x) F(x,Q,w) - u, range, extendInt = "yes")$root
    f <- function(y, Q, w, u = 1:99/100){
        z <- sapply(u, function(u) Finv(u, Q, w, range(Q)))
        approx(density(z), xout = y)$y
```

```
        }
    D <- function(x, y, Q, w, g)
        (g/(1-g)) * ((y <= x) -rowSums((Q <= x) * w))
    logLik <- function(eps, t, w, Do, Dq)
        -sum((1-t) * (eps * Do - log(rowSums(exp(eps * Dq) * w))))
    gradLik <- function(eps, t, w, Do, Dq)
        -sum((1-t) * (Do - rowSums(Dq * exp(eps * Dq) * w))/
            rowSums(exp(eps * Dq) * w))
    # Compute CQF of the control observations
    rqf <- rq(y ~ . - t, tau = taus, data = Z, subset = (t == 0))
    cqf <- predict(rqf, type = "Qhat", newdata = Z, stepfun = TRUE)
    cqf <- lapply(cqf, rearrange)
    Q <- t(sapply(cqf, function(f)f(unique(knots(f)))))
    y <- Z$y
    t <- Z$t
    # Compute propensity score
    fitT  <- glm(t ~ . - y, data = Z,  family = binomial)
    g <- predict(fitT, type = "response")
    # Initialization:
        n <- nrow(Q)
        m <- ncol(Q)
        w <- matrix(1/m, n, m)
        it <- 1
        crit <- TRUE
        q0 <- Finv(tau, Q[t == 1,], w[t == 1,], range(Q))
        q1 <- quantile(y[t == 1], tau)
        f0 <- f(q0, Q[t == 1,], w[t == 1])
        f1 <- approx(density(y[t == 1]), xout = q1)$y
        eif0 <- ((1-t) * D(q0, y, Q, w, g) +
            t * (rowSums((Q <= q0) * w) - tau))/(f0 * mean(t)) # (7)
        eif1 <-  t * ((y <= q1) - tau)/(f1 * mean(t))
        eif <- eif1 - eif0
        naive <- q1 - q0
        onestep <- q1 - q0 - mean(eif)
    while(crit && it <= maxit){
        Do <- D(q0, y, Q, w, g)
        Dq <- apply(Q, 2, function(y) D(q0, y, Q, w, g))
        eps <- optim(par = 0, logLik, gradLik, method = "BFGS",
                     t = t, w = w, Do = Do, Dq = Dq)$par
        w <- exp(eps * Dq) * w/rowSums(exp(eps * Dq) * w)
        q0 <- Finv(tau, Q[t == 1,], w[t == 1,], range(y))
        it <- it + 1
        crit <- (abs(eps) > 1e-3/n^(3/5))
        }
    f0 <- f(q0, Q[t == 1,], w[t == 1])
    eif0 <- ((1-t) * D(q0, y, Q, w, g) +
        t * (rowSums((Q <= q0) * w) - tau))/(f0 * mean(t)) # (7)
    eif <- eif1 - eif0
    firpo <- q1 - coef(rq(y ~ 1, tau = tau, data = Z, weights = (1-t) * g/(1-g)))
    list(tmle = q1 - q0, naive = naive, firpo = firpo, onestep = onestep,
         se = sqrt(var(eif)/n), nit = it)
}
```

In some respects the code for the QTET is simpler than the `tmlex` code we saw earlier. There is no fancy updating of the submodel function with `do.call`, etc., everything is dealt with explicitly in the way that the array `Dq` is updated. But essentially we have the same algorithmic structure as in earlier TMLE procedures:

Initialize p_Y
While(ϵ too big)
    Construct new parametric submodel
    Find MLE for ϵ in new model
    Update p_Y
Return

I wouldn't claim that the foregoing description will enable you to leap out of upper story windows with confidence, and even so you should always remember that Captain Hook and his pirates are lurking down below. But perhaps with some safety netting and further experimentation it will help to guide further exploration of the TMLE paradigm. I believe that the QTE extension of Díaz (2015) is a particularly promising development.

References

Bang H, Robins JM. 2005. Doubly robust estimation in missing data and causal inference models. *Biometrics* **61**: 962–973.

Chernozhukov V, Fernandez-Val I, Galichon A. 2009. Improving point and interval estimators of monotone functions by rearrangement. *Biometrika* : asp030.

Díaz I. 2015. Efficient estimation of quantiles in missing data models. *arXiv:1512.08110* .

Díaz I, Rosenblum M. 2015. Targeted maximum likelihood estimation using exponential families. *The international journal of biostatistics* **11**: 233–251.

Firpo S. 2007. Efficient semiparametric estimation of quantile treatment effects. *Econometrica* **75**: 259–276.

Horvitz DG, Thompson DJ. 1952. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association* **47**: 663–685.

Kang JD, Schafer JL. 2007. Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical science* : 523–539.

Koenker R, Yoon J. 2009. Parametric links for binary choice models: A Fisherian–Bayesian colloquy. *Journal of Econometrics* **152**: 120–130.

Marschak J. 1953. Economic measurements for policy and prediction. In Hood W, Koopmans T (eds.) *Studies in Econometric Method*. Yale, 1–26.

McCullagh P, Nelder JA. 1989. *Generalized linear models*. Chapman-Hall.

Rotnitzky A, Vansteelandt S. 2014. Double-robust methods. In *Handbook of missing data methodology*. CRC Press, 185–212.

van der Laan MJ, Rose S. 2011. *Targeted learning: causal inference for observational and experimental data*. Springer.

van der Laan MJ, Rubin D. 2006. Targeted maximum likelihood learning. *The International Journal of Biostatistics* **2**: 1–40.

van der Vaart AW. 2000. *Asymptotic statistics*. Cambridge University Press.