# A Ride in Targeted Learning Territory 

## Titre: Une virée en territoire d'apprentissage ciblé

## David Benkeser ${ }^{1}$ and Antoine Chambaz ${ }^{2}$


#### Abstract

Our ambition is to present a gentle introduction to the field of targeted learning. As an example, we consider statistical inference on a simple causal quantity that is ubiquitous in the causal literature. We use this exemplar parameter to introduce key concepts that can be applied to more complicated problems. The introduction weaves together two main threads, one theoretical and the other computational. It also contains exercises. The code is written in the programming language $R$, which is widely used among statisticians and data scientists to develop statistical software and data analysis. It uses tlrider, a package that we built specifically for this project.


Résumé : Notre ambition est de présenter une introduction accessible au domaine de l'apprentissage ciblé. À titre d'exemple, nous considérons l'inférence statistique d'une quantité causale simple souvent rencontrée dans la littérature causale. Nous utilisons ce paramètre exemplaire pour introduire des concepts-clefs qui jouent un rôle important dans des problèmes plus difficiles. L'introduction entrelace deux fils narratifs, l'un théorique et l'autre computationnel. Elle contient aussi des exercices. Le code est écrit dans le langage R, qui est largement utilisé par les statisticiens et data scientists pour développer des programmes statistiques et des analyses de données. Il s'appuie sur le paquetage tlrider, composé spécifiquement pour ce projet.

Keywords: causality, targeted learning
Mots-clés : apprentissage ciblé, causalité
AMS 2000 subject classifications: 62F12, 62F35, 62G08, 62G20

## 1. A ride

### 1.1. Introduction

Our ambition is to present a gentle introduction to the field of targeted learning. ${ }^{1}$ As an example, we consider statistical inference on a simple causal quantity that is ubiquitous in the causal literature. We use this exemplar parameter to introduce key concepts that can be applied to more complicated problems. The introduction weaves together two main threads, one theoretical and the other computational.

### 1.1.1. A causal story

We focus on a causal story where a random reward (a real number between 0 and 1 ) is given based on an action undertaken (one among two) and the random context where the action is performed

[^0](summarized by a real number between 0 and 1). The causal quantity of interest is the average difference of the two counterfactual rewards. This is a story as old as time. Should we take the red pill or the blue pill? Should we show our customers advertisement A or advertisement B? Should we require individuals to undergo cancer screening? At their core, each of these questions is asking what action should be taken to maximize a "reward."

We will build several estimators and discuss their respective merits, theoretically and computationally. The construction of the most involved estimator will unfold in targeted learning territory, at the frontier of machine learning and semiparametrics, the statistical theory of inference based on semiparametric models.

### 1.1.2. The tlrider package

The computational illustrations will be developed based on the companion package tlrider. The package can be installed by running the following code:

```
devtools::install_github("achambaz/tlride/tlrider")
```

The version used in this document is 1.1.0.
Additional packages are also required, including tidyverse (Wickham and Grolemund, 2016), caret (Kuhn, 2020) and ggdag (Barrett, 2018). Assuming that these are installed too, we can run the next chunk of code:

```
set.seed(3141516) ## because reproducibility matters...
library(tidyverse)
library(caret)
library(ggdag)
library(tlrider)
```


### 1.1.3. What we will discuss

To begin, we discuss the nature of the parameter of interest, viewing it as the value of a statistical mapping evaluated at the law of the data (Section 2), with an emphasis on the smoothness and double-robustness properties inherited from the mapping (Sections 3 and 4). We then turn to the estimation of the parameter of interest. We first introduce and comment upon a simple inference strategy assuming provisionally that a relevant feature of the law of the data is known to us (Section 6). Second, we present the notion of nuisance parameters and adopt an algorithmic stance on their estimation (Section 7). Third, we introduce and comment upon two "naive" inference strategies (Section 8), the one-step correction procedure (Section 9) and, finally, the targeted minimum loss estimation procedure tailored to the inference of the parameter of main interest. In the appendix, we collect our notation (Section A), and present some results that are used in the main text and their proofs (Sections B and C).

### 1.2. A simulation study

### 1.2.1. Reproducible experiment as a law

We are interested in a reproducible experiment. Every time this experiment is run, it generates an observation that we call $O$. We view $O$ as a random variable drawn from the law of the experiment that we denote by $P_{0}$.

We view $P_{0}$ as an element of the model $\mathscr{M}$. The model is a collection of laws. In particular, the model contains all laws that we think may plausibly describe the law of the experiment. Thus, the choice of model is based on our scientific knowledge of the experiment. The more we know about the experiment, the smaller is $\mathscr{M}$. In all our examples, we use large models that reflect a lack of knowledge about many aspects of the experiment.

### 1.2.2. A synthetic reproducible experiment

Instead of considering a real-life reproducible experiment, we focus for pedagogical purposes on a synthetic reproducible experiment. Thus we can from now on take on two different roles: that of an oracle knowing completely the nature of the experiment, and that of a statistician eager to know more about the experiment by observing some of its outputs.

Let us run the example built into the tlrider package:

```
example(tlrider)
```

A few objects have been defined:

```
ls()
#> [1] "another_experiment" "experiment" "expit"
#> [4] "filter" "logit" "sigmaO"
```

The function expit implements the link function expit : $\mathbb{R} \rightarrow] 0,1[$ given by expit $(x) \doteq(1+$ $\left.e^{-x}\right)^{-1}$. The function logit implements its inverse function $\left.\operatorname{logit}:\right] 0,1[\rightarrow \mathbb{R}$ given by $\operatorname{logit}(p) \doteq$ $\log [p /(1-p)]$.

Let us take a look at experiment:

```
experiment
#> A law for (W,A,Y) in [0,1] x {0,1} x [0,1].
#>
#> If the law is fully characterized, you can use method
#> 'sample_from' to sample from it.
#>
#> If you built the law, or if you are an _oracle_, you can also
#> use methods 'reveal' to reveal its relevant features (QW, Gbar,
#> Qbar, qY -- see '?reveal'), and 'alter' to change some of them.
#>
#> If all its relevant features are characterized, you can use
#> methods 'evaluate_psi' to obtain the value of 'Psi' at this law
#> (see '?evaluate_psi') and 'evaluate_eic' to obtain the efficient
#> influence curve of 'Psi' at this law (see '?evaluate_eic').
```

The law $P_{0}$ of the synthetic experiment experiment built by us generates a generic observation $O$ that decomposes as

$$
O \doteq(W, A, Y) \in[0,1] \times\{0,1\} \times[0,1]
$$

We interpret $W$ as a real valued summary measure of a random context where an action $A$ chosen among two is undertaken, leading to a real valued reward $Y$.

We can sample from the experiment (simply run ?sample_from to see the man page of method sample_from). The next chunk of code runs the experiment five times, independently:

```
(five_obs <- sample_from(experiment, n = 5))
#> W A Y
#> [1,] 0.414 1 0.996
#> [2,] 0.409 1 0.669
#> [3,] 0.404 0 0.825
#> [4,] 0.462 1 0.539
#> [5,] 0.404 1 0.986
```


### 1.2.3. Revealing experiment

Acting as oracles, we can peek into experiment and reveal a selection of relevant features (simply run ?reveal to see the man page of method reveal). Made by us, the selection exhibits features that will play an important role in the text.

```
relevant_features <- reveal(experiment)
names(relevant_features)
\#> [1] "QW" "Gbar" "Qbar" "qY" "sample_from"
```

We have an oracular knowledge of experiment and can thus comment upon the features of $P_{0}$ revealed in relevant_features.

QW The QW feature describes the marginal law of $W$, that we call $Q_{0, W}{ }^{2}$

```
relevant_features$QW
#> function(W,
#> mixture_weights = c(1/10, 9/10, 0),
#> mins = c(0, 11/30, 0),
#> maxs = c(1, 14/30, 1)) {
#> out <- sapply(1:length(mixture_weights),
#> function(ii){
#> mixture_weights[ii] *
#> stats::dunif(W,
#> min = mins[ii],
#> max = maxs[ii])
#> })
#> return(rowSums (out))
```

[^1]```
#> }
#> <environment: Oxf3ca190>
```

It appears that $Q_{0, W}$ is a mixture of the uniform laws over $[0,1]$ (weight $1 / 10$ ) and $[11 / 30,14 / 30]$ (weight $9 / 10$ ). ${ }^{3}$

Gbar The Gbar feature describes the conditional probability of action $A=1$ given $W$. For each $a \in\{0,1\}$, we denote

$$
\begin{aligned}
\bar{G}_{0}(W) & \doteq \operatorname{Pr}_{P_{0}}(A=1 \mid W), \\
\ell \bar{G}_{0}(a, W) & \doteq \operatorname{Pr}_{P_{0}}(A=a \mid W) .
\end{aligned}
$$

Obviously,

$$
\ell \bar{G}_{0}(A, W)=A \bar{G}_{0}(W)+(1-A)\left(1-\bar{G}_{0}(W)\right) .
$$

```
relevant_features$Gbar
#> function(W) {
#> expit(1 + 2*W-4* sqrt(abs((W - 5/12))))
#> }
#> <environment: Oxf3ca190>
```

Note how real numbers of the form $1+2 W-4 * \sqrt{|W-5 / 12|}$ ) are mapped into the interval $[0,1]$ by the expit link function. We refer the reader to Figure 6 for a visualization of $\bar{G}_{0}$.
qY The qY feature describes the conditional density of $Y$ given $A$ and $W$. For each $y \in] 0,1[$, we denote by $q_{0, Y}(y, A, W)$ the conditional density evaluated at $y$ of $Y$ given $A$ and $W$.

```
relevant_features$qY
#> function(obs, Qbar, shape10 = 2, shape11 = 3){
#> A <- obs[, "A"]
#> AW <- obs[, c("A", "W")]
#> QAW <- Qbar(AW)
#> shape1 <- ifelse(A == 0, shape10, shape11)
#> stats::dbeta(Y,
#> shape1 = shape1,
#> shape2 = shape1 * (1 - QAW) / QAW)
#> }
> <environment: Oxf3ca190>
```

It appears that the conditional law of $Y$ given $A$ and $W$ is the Beta law with conditional mean and variance characterized by the Qbar feature of experiment (see below) and the shape 10 and shape11 parameters.

[^2]Qbar As for the Qbar feature, it describes the conditional mean of $Y$ given $A$ and $W$.

```
relevant_features$Qbar
#> function(AW) {
#> A <- AW[, "A"]
#> W <- AW[, "W"]
#> A * (cos((-1/2 + W) * pi) * 2/5 + 1/5 +
#> (1/3 <= W छ W <= 1/2) / 5 +
#> (W >= 3/4)*(W-3/4)*2) +
#> (1-A)*(sin(4*W`2*pi)/4 + 1/2)
#> }
#> <bytecode: 0x104fa668>
#> <environment: Oxf3ca190>
```

We denote $\bar{Q}_{0}(A, W)=\mathrm{E}_{P_{0}}(Y \mid A, W)$ the conditional mean of $Y$ given $A$ and $W$. Note how $\bar{Q}_{0}(A, W)$ does depend heavily on $A$ and $W$. We refer the reader to Section 1.3 for a visualization of $\bar{Q}_{0}$.
sample_from Finally, the sample_from feature is the function called by method sample_from when it is applied to an object of class LAW, like experiment.

```
relevant_features$sample_from
#> function(n, ideal = FALSE) {
#> ## preliminary
#> n<- R.utils::Arguments$getInteger(n, c(1, Inf))
#> ideal <- R.utils::Arguments$getLogical(ideal)
#> ## ## 'Gbar' and 'Qbar' factors
#> Gbar <- experiment$.Gbar
#> Qbar <- experiment$.Qbar
#> ## sampling
#> ## ## context
#> params <- formals(experiment$.QW)
#> mixture_weights <- eval(params$mixture_weights)
#> mins <- eval(params$mins)
#> maxs <- eval(params$maxs)
#> W <- sample_from_mixture_of_uniforms(n, mixture_weights,
#> mins, maxs)
#> ## ## counterfactual rewards
#> zeroW <- cbind}(A=0,W
#> oneW <- cbind(A = 1, W)
#> Qbar_zeroW <- Qbar(zeroW)
#> Qbar_oneW <- Qbar(oneW)
#> Yzero <- stats::rbeta(n,
#> shape1 = 2,
#> shape2 = 2 * (1 - Qbar_zeroW) / Qbar_zeroW)
#> Yone <- stats::rbeta(n,
#> shape1 = 3,
#> shape2 = 3 * (1 - Qbar_oneW) / Qbar_oneW)
#> ## ## action undertaken
```

```
#> A <- stats::rbinom(n, size = 1, prob = Gbar(W))
#> ## ## actual reward
#> Y<- A * Yone + (1 - A) * Yzero
#> ## ## observation
#> if (ideal) {
#> obs <- cbind(W = W, Yzero = Yzero, Yone = Yone, A = A, Y=Y)
#> } else {
#> obs <- cbind}(W=W,A=A,Y=Y
#> }
#> return(obs)
#> }
#> <bytecode: Oxedf6bc0>
> <environment: Oxf3ca190>
```

We will comment upon the ideal argument in the above sample_from feature in Section 2.1.

### 1.3. Visualization

1. Run the following chunk of code. It visualizes the conditional mean $\bar{Q}_{0}$.
```
Gbar <- relevant_features$Gbar
Qbar <- relevant_features$Qbar
QW <- relevant_features$QW
features <- tibble(w = seq(0, 1, length.out = 1e3)) %>%
    mutate(Qw = QW(w),
            Gw = Gbar(w),
            Q1w = Qbar(cbind(A = 1, W = w)),
            QOw = Qbar(cbind(A = 0, W = w)),
            blip_Qw = Q1w - Q0w)
features %>% select(-Qw, -Gw) %>%
    rename("Q(1,.)" = Q1w,
            "Q(0,.)" = Q0w,
            "Q(1,.) - Q(0,.)" = blip_Qw) %>%
    pivot_longer(-w, names_to = "f", values_to = "value") %>%
    ggplot() +
    geom_line(aes(x = w, y = value, color = f), size = 1) +
    labs(y = "f(w)", title = bquote("Visualizing" ~ bar(Q)[0])) +
    ylim(NA, 1)
```


2. Adapt the above chunk of code to visualize the marginal density $Q_{0, W}$ and conditional probability $\bar{G}_{0}$.

### 1.4. Make your own experiment

You can easily make your own experiment.

1. Check out the man page of method alter by running ?alter.
2. Run the following chunk of code:
```
my_experiment <- LAW() ## creates an object of class 'LAW'
alter(my_experiment, ## characterize its relevant features
    QW = function(W) {
        out <- rep_len(0, length(W))
        out[W == 0] <- 1/4
        out[W == 1] <- 3/4
        return(out)
    },
    Gbar = function(W) {
        out <- rep_len(0, length(W))
        out[W == 0] <- 1/3
        out[W == 1] <- 3/5
        return(out)
    },
    Qbar = function(AW) {
        probs <- matrix(c(1/2, 2/3, 7/8, 4/5), ncol = 2,
                        dimnames = list(c("A=0", "A=1"),
                                    c("W=0", "W=1")))
        probs[cbind(AW[, "A"] + 1, AW[, "W"] + 1)]
    },
    qY = function(obs) {
        probs <- matrix(c(1/2, 2/3, 7/8, 4/5), ncol = 2,
                        dimnames = list(c("A=0", "A=1"),
```

```
                                    c("W=0", "W=1")))
    probs <- probs[cbind(obs[, "A"] + 1, obs[, "W"] + 1)]
    obs[, "Y"] * probs + (1 - obs[, "Y"]) * (1 - probs)
},
sample_from = function(n) {
    ## preliminary
    n <- R.utils::Arguments$getInteger(n, c(1, Inf))
    ## 'QW', 'Gbar' and 'Qbar' features
    QW <- my_experiment$.QW
    Gbar <- my_experiment$.Gbar
    Qbar <- my_experiment$.Qbar
    ## sampling
    W <- rbinom(n, size = 1, prob = QW(1))
    A <- rbinom(n, size = 1, prob = Gbar(W))
    AW <- cbind(W = W, A = A)
    Y <- rbinom(n, size = 1, Qbar(AW))
    return(cbind(AW, Y = Y))
})
```

3. What does the next chunk do?
```
(sample_from(my_experiment, 3))
#> W A Y
#> [1,] 1 0 1
#> [2,] 0 0 1
#> [3,] 1 1 1
```

4. Characterize entirely the law of my_experiment. Hint:
```
obs <- sample_from(my_experiment, 1e4)
obs %>% as_tibble %>% group_by(W, A, Y) %>%
    summarize(how_many = n()) %>% ungroup
#> # A tibble: 8 x 4
#> W A Y how_many
#> <int> <int> <int> <int>
#> 1 0 0 0 826
#> 2 0 0 1 1 833
#> 3 0 1 0 286
#> 4 0
#> 5 1 0
#> 6 1 0 0 1 
#> # ... with 2 more rows
obs %>% as_tibble %>% group_by(W, A) %>%
    summarize(prob = mean(Y)) %>% ungroup
#> # A tibble: 4 x 3
#> W A prob
#> <int> <int> <dbl>
#> 1 0 0 0.502
#> 2 0 1 0.662
```

```
#> 3 1 1 0 0.872
#> 4 1 1 0.798
```

5. Now, make your own experiment.

## 2. The parameter of interest

### 2.1. The parameter of interest

### 2.1.1. Definition

It happens that we especially care for a finite-dimensional feature of $P_{0}$ that we denote by $\psi_{0}$. Its definition involves two of the aforementioned infinite-dimensional features, the marginal law $Q_{0, W}$ of $W$ and the conditional mean $\bar{Q}_{0}$ of $Y$ given $A$ and $W$ :

$$
\begin{align*}
\psi_{0} & \doteq \int\left(\bar{Q}_{0}(1, w)-\bar{Q}_{0}(0, w)\right) d Q_{0, W}(w)  \tag{1}\\
& =\mathrm{E}_{P_{0}}\left(\mathrm{E}_{P_{0}}(Y \mid A=1, W)-\mathrm{E}_{P_{0}}(Y \mid A=0, W)\right)
\end{align*}
$$

Acting as oracles, we can compute explicitly the numerical value of $\psi_{0}$. The evaluate_psi method makes it very easy (simply run ?estimate_psi to see the man page of the method):
(psi_zero <- evaluate_psi(experiment))
\#> [1] 0.0832

### 2.1.2. A causal interpretation

Our interest in $\psi_{0}$ is of causal nature. Taking a closer look at the sample_from feature of experiment reveals indeed that the random making of an observation $O$ drawn from $P_{0}$ can be summarized by the following directed acyclic graph:

```
dagify(
    Y ~ A + Y1 + Y0, A ~ W, Y1 ~ W, Y0 ~ W,
    labels = c(Y = "Actual reward",
                A = "Action",
                Y1 = "Counterfactual reward\n of action 1",
                        YO = "Counterfactual reward\n of action 0",
                        W = "Context of action"),
    coords = list(
        x = c(W = 0, A = -1, Y1 = 1.5, YO = 0.25, Y = 1),
        y = c(W = 0, A = -1, Y1 = -0.5, Y0 = -0.5, Y = -1)),
    outcome = "Y",
    exposure = "A",
    latent = c("Y0", "Y1")) %>% tidy_dagitty %>%
    ggdag(text = TRUE, use_labels = "label") + theme_dag_grey()
```

In words, the experiment unfolds like this (see also Section B.1):


Figure 1. Directed acyclic graph summarizing the inner causal mechanism at play in experiment.

1. a context of action $W \in[0,1]$ is randomly generated;
2. two counterfactual rewards $Y_{0} \in[0,1]$ and $Y_{1} \in[0,1]$ are generated conditionally on $W$;
3. an action $A \in\{0,1\}$ (among two possible actions called $a=0$ and $a=1$ ) is undertaken, (i) knowing the context but not the counterfactual rewards, and (ii) in such a way that both actions can always be considered;
4. the action yields a reward $Y$, which equals either $Y_{0}$ or $Y_{1}$ depending on whether action $a=0$ or $a=1$ has been undertaken;
5. summarize the course of the experiment with $O \doteq(W, A, Y)$, thus concealing $Y_{0}$ and $Y_{1}$.

The above description of the experiment is useful to reinforce what it means to run the "ideal" experiment by setting argument ideal to TRUE in a call to sample_from for experiment (see Section 2.1.3). Doing so triggers a modification of the nature of the experiment, enforcing that the counterfactual rewards $Y_{0}$ and $Y_{1}$ be part of the summary of the experiment eventually. In light of the above enumeration,

$$
\mathbb{O} \doteq\left(W, Y_{0}, Y_{1}, A, Y\right)
$$

is output, as opposed to its summary measure $O$. This defines another experiment and its law, that we denote $\mathbb{P}_{0}$.

It is straightforward to show that

$$
\begin{align*}
\psi_{0} & =\mathrm{E}_{\mathbb{P}_{0}}\left(Y_{1}-Y_{0}\right)  \tag{2}\\
& =\mathrm{E}_{\mathbb{P}_{0}}\left(Y_{1}\right)-\mathrm{E}_{\mathbb{P}_{0}}\left(Y_{0}\right)
\end{align*}
$$

Thus, $\psi_{0}$ describes the average difference of the two counterfactual rewards. In other words, $\psi_{0}$ quantifies the difference in average of the reward one would get in a world where one would always enforce action $a=1$ with the reward one would get in a world where one would always enforce action $a=0$. This said, it is worth emphasizing that $\psi_{0}$ is a well-defined parameter beyond its causal interpretation, and that it describes a standardized association between the action $A$ and reward $Y$.

### 2.1.3. A causal computation

We can use our position as oracles to sample observations from the ideal experiment. We call sample_from for experiment with its argument ideal set to TRUE in order to numerically approximate $\psi_{0}$. By the law of large numbers, the following code approximates $\psi_{0}$ and shows it approximate value.

```
B <- 1e6
ideal_obs <- sample_from(experiment, B, ideal = TRUE)
(psi_approx <- mean(ideal_obs[, "Yone"] - ideal_obs[, "Yzero"]))
#> [1] 0.083
```

The object psi_approx contains an approximation to $\psi_{0}$ based on B observations from the ideal experiment. The random sampling of observations results in uncertainty in the numerical approximation of $\psi_{0}$. This uncertainty can be quantified by constructing a $95 \%$ confidence interval for $\psi_{0}$. The central limit theorem and Slutsky's lemma allow us to build such an interval as follows.

```
sd_approx <- sd(ideal_obs[, "Yone"] - ideal_obs[, "Yzero"])
alpha <- 0.05
(psi_approx_CI <- psi_approx + c(-1, 1) *
    qnorm(1 - alpha / 2) * sd_approx / sqrt(B))
#> [1] 0.0824 0.0836
```

We note that the interpretation of this confidence interval is that in $95 \%$ of draws of size $B$ from the ideal data generating experiment, the true value of $\psi_{0}$ will be contained in the generated confidence interval.

### 2.2. An alternative parameter of interest

Equality (2) shows that parameter $\psi_{0}$ (1) is the difference in average rewards if we enforce action $a=1$ rather than $a=0$. An alternative way to describe the rewards under different actions involves quantiles as opposed to averages.

Let

$$
Q_{0, Y}(y, A, W) \doteq \int_{0}^{y} q_{0, Y}(u, A, W) d u
$$

be the conditional cumulative distribution of reward $Y$ given $A$ and $W$, evaluated at $y \in] 0,1[$, that is implied by $P_{0}$. For each action $a \in\{0,1\}$ and $\left.c \in\right] 0,1[$, introduce

$$
\begin{equation*}
\gamma_{0, a, c} \doteq \inf \{y \in] 0,1\left[: \int Q_{0, Y}(y, a, w) d Q_{0, W}(w) \geq c\right\} \tag{3}
\end{equation*}
$$

(Note: inf merely generalizes min, accounting for the fact that the minimum may fail to be achieved.)

It is not very difficult to check (see Problem 1 below) that

$$
\begin{equation*}
\gamma_{0, a, c}=\inf \{y \in] 0,1\left[: \operatorname{Pr}_{\mathbb{P}_{0}}\left(Y_{a} \leq y\right) \geq c\right\} \tag{4}
\end{equation*}
$$

Thus, $\gamma_{0, a, c}$ can be interpreted as the $c$-th quantile reward when action $a$ is enforced. The difference

$$
\begin{equation*}
\delta_{0, c} \doteq \gamma_{0,1, c}-\gamma_{0,0, c} \tag{5}
\end{equation*}
$$

is the $c$-th quantile counterpart to parameter $\psi_{0}(1)$.

1. 2 Prove (4).
2. 2 Compute the numerical value of $\gamma_{0, a, c}$ for each $(a, c) \in\{0,1\} \times\{1 / 4,1 / 2,3 / 4\}$ using the appropriate features of experiment (see relevant_features). Based on these results, report the numerical value of $\delta_{0, c}$ for each $c \in\{1 / 4,1 / 2,3 / 4\}$.
3. Approximate the numerical values of $\gamma_{0, a, c}$ for each $(a, c) \in\{0,1\} \times\{1 / 4,1 / 2,3 / 4\}$ by drawing a large sample from the "ideal" data experiment and using empirical quantile estimates. Deduce from these results a numerical approximation to $\delta_{0, c}$ for $c \in\{1 / 4,1 / 2,3 / 4\}$. Confirm that your results closely match those obtained in the previous problem.

### 2.3. The statistical mapping of interest

The noble way to define a statistical parameter is to view it as the value of a statistical mapping at the law of the experiment of interest. Beyond the elegance, this has paramount statistical implications.

### 2.3.1. Opening discussion

Oftentimes, the premise of a statistical analysis is presented like this. One assumes that the law $P_{0}$ of the experiment of interest belongs to a statistical model

$$
\left\{P_{\theta}: \theta \in T\right\}
$$

(where $T$ is some index set). The statistical model is identifiable, meaning that if two elements $P_{\theta}$ and $P_{\theta^{\prime}}$ coincide, then necessarily $\theta=\theta^{\prime}$. Therefore, there exists a unique $\theta_{0} \in T$ such that $P_{0}=P_{\theta_{0}}$, and one wishes to estimate $\theta_{0}$.

For instance, each $P_{\theta}$ could be the Gaussian law with mean $\theta \in T \doteq \mathbb{R}$ and variance 1 , and one could wish to estimate the mean $\theta_{0}$ of $P_{0}$. To do so, one could rely on $n$ observations $X_{1}, \ldots$, $X_{n}$ drawn independently from $P_{0}$. The empirical mean

$$
\theta_{n} \doteq \frac{1}{n} \sum_{i=1}^{n} X_{i}
$$

estimates $\theta_{0}$. More generally, if we assume that $\operatorname{Var}_{P_{0}}\left(X_{1}\right)$ is finite, then $\theta_{n}$ satisfies many useful properties. In particular, it can be used to construct confidence intervals.

Of course, the mean of a law is defined beyond the small model $\left\{P_{\theta}: \theta \in \mathbb{R}\right\}$. Let $\mathscr{M}$ be the set of laws $P$ on $\mathbb{R}$ such that $\operatorname{Var}_{P}(X)$ is finite. In particular, $P_{0} \in \mathscr{M}$. For every $P \in \mathscr{M}$, the mean $\mathrm{E}_{P}(X)$ is well defined. Thus, we can introduce the statistical mapping $\Theta: \mathscr{M} \rightarrow \mathbb{R}$ given by

$$
\Theta(P) \doteq \mathrm{E}_{P}(X)
$$

Interestingly, the empirical measure $P_{n}{ }^{4}$ is an element of $\mathscr{M}$. Therefore, the statistical mapping $\Theta$ can be evaluated at $P_{n}$ :

$$
\Theta\left(P_{n}\right)=\frac{1}{n} \sum_{i=1}^{n} X_{i}=\theta_{n}
$$

We recover the empirical mean, and understand that it is a substitution estimator of the mean: in order to estimate $\Theta\left(P_{0}\right)$, we substitute $P_{n}$ for $P_{0}$ within $\Theta$. ${ }^{5}$

Substitution-based estimators are particularly valuable notably because they, by construction, satisfy all the constraints to which the targeted parameter is subjected. For example, if $X$ is a binary random variable and the support of all distributions in our model is $\{0,1\}$, then $\Theta$ can be interpreted as the probability that $X=1$, a quantity known to live in the interval $[0,1]$. A substitution estimator will also be guaranteed to fall into this interval. Some of the estimators that we will build together are substitution-based, some are not.

### 2.3.2. The parameter as the value of a statistical mapping at the experiment

We now go back to our main topic of interest. Suppose we know beforehand that $O$ drawn from $P_{0}$ takes its values in $\mathscr{O} \doteq[0,1] \times\{0,1\} \times[0,1]$ and that $\bar{G}_{0}(W) \doteq \operatorname{Pr}_{P_{0}}(A=1 \mid W)$ is bounded away from zero and one $Q_{0, W}$-almost surely (this is the case indeed). Then we can define model $\mathscr{M}$ as the set of all laws $P$ on $\mathscr{O}$ such that

$$
\bar{G}(W) \doteq \operatorname{Pr}_{P}(A=1 \mid W)
$$

is bounded away from zero and one $Q_{W}$-almost surely, where $Q_{W}$ is the marginal law of $W$ under $P$.

Let us also define generically $\bar{Q}$ as

$$
\bar{Q}(A, W) \doteq \mathrm{E}_{P}(Y \mid A, W)
$$

Note how we have suppressed the dependence of $\bar{G}$ and $\bar{Q}$ on $P$ for notational simplicity.
Central to our approach is viewing $\psi_{0}$ as the value at $P_{0}$ of the statistical mapping $\Psi$ from $\mathscr{M}$ to $[0,1]$ characterized by

$$
\begin{align*}
\Psi(P) & \doteq \int\left(\bar{Q}_{P}(1, w)-\bar{Q}_{P}(0, w)\right) d Q_{W}(w)  \tag{6}\\
& =\mathrm{E}_{P}\left(\bar{Q}_{P}(1, W)-\bar{Q}_{P}(0, W)\right),
\end{align*}
$$

a clear extension of (1) where, for once, we make the dependence of $\bar{Q}$ on $P$ explicit to emphasize how $\Psi(P)$ truly depends on $P$.

[^3]
### 2.3.3. The value of the statistical mapping at another experiment

When we ran example(tlrider) earlier, we created an object called another_experiment:

```
another_experiment
#> A law for (W,A,Y) in [0,1] x {0,1} x [0,1].
#>
#> If the law is fully characterized, you can use method
#> 'sample_from' to sample from it.
#>
#> If you built the law, or if you are an _oracle_, you can also
#> use methods 'reveal' to reveal its relevant features (QW, Gbar,
#> Qbar, qY -- see '?reveal'), and 'alter' to change some of them.
#>
#> If all its relevant features are characterized, you can use
#> methods 'evaluate_psi' to obtain the value of 'Psi' at this law
#> (see '?evaluate_psi') and 'evaluate_eic' to obtain the efficient
#> influence curve of 'Psi' at this law (see '?evaluate_eic').
reveal(another_experiment)
#> $QW
#> function(x, min = 1/10, max = 9/10){
#> stats::dunif(x, min = min, max = max)
#> }
<environment: Ox10837e90>
$Gbar
function(W) {
            sin((1 + W)*pi / 6)
            }
    <environment: Ox10837e90>
    $Qbar
    function(AW,h) {
        A <- AW[, "A"]
        W <- AW[, "W"]
            expit( logit(A * W + (1 - A) * W`2 ) +
                h * 10 * sqrt(W) * A )
        }
    <environment: Ox10837e90>
    $qY
    function(obs, Qbar, shape1 = 4){
            AW <- obs[, c("A", "W")]
            QAW <- Qbar(AW)
            stats::gdbeta(Y,
                                shape1 = shape1,
                                shape2 = shape1 * (1 - QAW) / QAW)
            }
    <environment: Ox10837e90>
```

```
$sample_from
function(n, h) {
            ## preliminary
            n<- R.utils::Arguments$getInteger(n, c(1, Inf))
            h <- R.utils::Arguments$getNumeric(h)
            ## ## 'Gbar' and 'Qbar' factors
            Gbar <- another_experiment$.Gbar
            Qbar <- another_experiment$.Qbar
            ## sampling
            ## ## context
            params <- formals(another_experiment$.QW)
            W <- stats::runif(n, min = eval(params$min),
                max = eval(params$max))
            ## ## action undertaken
            A <- stats::rbinom(n, size = 1, prob = Gbar(W))
            ## ## reward
            params <- formals(another_experiment$.qY)
            shape1 <- eval(params$shape1)
            QAW <- Qbar(cbind(A = A, W = W), h = h)
            Y <- stats::rbeta(n,
                                    shape1 = shape1,
                            shape2 = shape1 * (1 - QAW)/QAW)
        ## ## observation
        obs <- cbind(W = W, A = A, Y = Y)
        return(obs)
    }
<environment: Ox10837e90>
(two_obs_another_experiment <- sample_from(another_experiment, 2, h = 0))
#> W A Y
#> [1,] 0.101 0 0.00841
#> [2,] 0.620 1 0.51166
```

By taking an oracular look at the output of reveal (another_experiment), we discover that the law $\Pi_{0} \in \mathscr{M}$ encoded by default (i.e., with $\mathrm{h}=0$ ) in another_experiment differs starkly from $P_{0}$.

However, the parameter $\Psi\left(\Pi_{0}\right)$ is well defined. Straightforward algebra shows that $\Psi\left(\Pi_{0}\right)=$ $59 / 300$. The numeric computation below confirms the equality.

```
(psi_Pi_zero <- evaluate_psi(another_experiment, h = 0))
#> [1] 0.197
round(59/300, 3)
#> [1] 0.197
```


### 2.4. Alternative statistical mapping

We now resume the exercise of Section 2.2. Like we did in Section 2.3, we introduce a generic version of the relevant features $q_{0, Y}$ and $Q_{0, Y}$. Specifically, we define $q_{Y}(y, A, W)$ to be the conditional density of $Y$ given $A$ and $W$, evaluated at $y$, that is implied by a generic $P \in \mathscr{M}$. Similarly,
we use $Q_{Y}$ to denote the corresponding cumulative distribution function.
The covariate-adjusted $c$-th quantile reward for action $a \in\{0,1\}, \gamma_{0, a, c}(3)$, may be viewed as the value at $P_{0}$ of a mapping $\Gamma_{a, c}$ from $\mathscr{M}$ to $[0,1]$ characterized by

$$
\Gamma_{a, c}(P)=\inf \{y \in] 0,1\left[: \int Q_{Y}(y, a, w) d Q_{W}(w) \geq c\right\}
$$

The difference in $c$-th quantile rewards, $\delta_{0, c}$ (5), may similarly be viewed as the value at $P_{0}$ of a mapping $\Delta_{c}$ from $\mathscr{M}$ to $[0,1]$, characterized by

$$
\Delta_{c}(P) \doteq \Gamma_{1, c}(P)-\Gamma_{0, c}(P) .
$$

1. Compute the numerical value of $\Gamma_{a, c}\left(\Pi_{0}\right)$ for $(a, c) \in\{0,1\} \times\{1 / 4,1 / 2,3 / 4\}$ using the relevant features of another_experiment. Based on these results, report the numerical value of $\Delta_{c}\left(\Pi_{0}\right)$ for each $c \in\{1 / 4,1 / 2,3 / 4\}$.
2. Approximate the value of $\Gamma_{0, a, c}\left(\Pi_{0}\right)$ for $(a, c) \in\{0,1\} \times\{1 / 4,1 / 2,3 / 4\}$ by drawing a large sample from the "ideal" data experiment and using empirical quantile estimates. Deduce from these results a numerical approximation to $\Delta_{0, c}\left(\Pi_{0}\right)$ for each $c \in\{1 / 4,1 / 2,3 / 4\}$. Confirm that your results closely match those obtained in the previous problem.
3. Building upon the code you wrote to solve the previous problem, construct a confidence interval with asymptotic level $95 \%$ for $\Delta_{0, c}\left(\Pi_{0}\right)$, with $c \in\{1 / 4,1 / 2,3 / 4\}$.

### 2.5. Representations

In Section 2.3, we reoriented our view of the target parameter to be that of a statistical functional of the law of the observed data. Specifically, we viewed the parameter as a function of specific features of the observed data law, namely $Q_{W}$ and $\bar{Q}$.

### 2.5.1. Yet another representation

It is straightforward to show an equivalent representation of the parameter as

$$
\begin{align*}
\psi_{0} & =\int \frac{2 a-1}{\ell \bar{G}_{0}(a, w)} y d P_{0}(w, a, y) \\
& =\mathrm{E}_{P_{0}}\left(\frac{2 A-1}{\ell \bar{G}_{0}(A, W)} Y\right) . \tag{7}
\end{align*}
$$

Viewing again the parameter as a statistical mapping from $\mathscr{M}$ to $[0,1]$, it also holds that

$$
\begin{align*}
\Psi(P) & =\int \frac{2 a-1}{\ell \bar{G}(a, w)} y d P(w, a, y) \\
& =\mathrm{E}_{P}\left(\frac{2 A-1}{\ell \bar{G}_{0}(A, W)} Y\right) . \tag{8}
\end{align*}
$$

### 2.5.2. From representations to estimation strategies

Our reason for introducing this alternative view of the target parameter will become clear when we discuss estimation of the target parameter. Specifically, the representations (1) and (7) naturally suggest different estimation strategies for $\psi_{0}$, as hinted in Section 2.3.1. The former suggests building an estimator of $\psi_{0}$ using estimators of $\bar{Q}_{0}$ and of $Q_{W, 0}$. The latter suggests building an estimator of $\psi_{0}$ using estimators of $\ell \bar{G}_{0}$ and of $P_{0}$.

We return to these ideas in later sections.

### 2.6. Alternative representation

1. 2 Show that for $a^{\prime}=0,1, \gamma_{0, a^{\prime}, c}$ as defined in (3) can be equivalently expressed as

$$
\inf \{z \in] 0,1\left[: \int \frac{\mathbf{1}\left\{a=a^{\prime}\right\}}{\ell \overline{\boldsymbol{G}}\left(a^{\prime}, W\right)} \mathbf{1}\{y \leq z\} d P_{0}(w, a, y) \geq c\right\}
$$

## 3. Smoothness

### 3.1. Fluctuating smoothly

Within our view of the target parameter as a statistical mapping evaluated at the law of the experiment, it is natural to inquire of properties this functional enjoys. For example, we may be interested in asking how the value of $\Psi(P)$ changes as we consider laws that get nearer to $P$ in $\mathscr{M}$. If small deviations from $P_{0}$ result in large changes in $\Psi\left(P_{0}\right)$, then we might hypothesize that it will be difficult to produce stable estimators of $\psi_{0}$. Fortunately, this turns out not to be the case for the mapping $\Psi$, and so we say that $\Psi$ is a smooth statistical mapping.
To discuss how $\Psi(P)$ changes for distributions that get nearer to $P$ in the model, we require a more concrete notion of what it means to get near to a distribution in a model. The notion hinges on fluctuations (or fluctuating models).

### 3.1.1. The another_experiment fluctuation

In Section 2.3.3, we discussed the nature of the object called another_experiment that was created when we ran example(tlrider):

```
another_experiment
#> A law for (W,A,Y) in [0,1] x {0,1} x [0,1].
#>
#> If the law is fully characterized, you can use method
#> 'sample_from' to sample from it.
#>
#> If you built the law, or if you are an _oracle_, you can also
#> use methods 'reveal' to reveal its relevant features (QW, Gbar,
#> Qbar, qY -- see '?reveal'), and 'alter' to change some of them.
#>
#> If all its relevant features are characterized, you can use
#> methods 'evaluate_psi' to obtain the value of 'Psi' at this law
```

```
#> (see '?evaluate_psi') and 'evaluate_eic' to obtain the efficient
#> influence curve of 'Psi' at this law (see '?evaluate_eic').
```

The message is a little misleading. Indeed, another_experiment is not $a$ law but, rather, a collection of laws indexed by a real-valued parameter h. This oracular statement (we built the object!) is evident when one looks again at the sample_from feature of another_experiment:

```
reveal(another_experiment)$sample_from
    function(n, h) {
            ## preliminary
            n <- R.utils::Arguments$getInteger(n, c(1, Inf))
            h <- R.utils::Arguments$getNumeric(h)
                        ## ## 'Gbar' and 'Qbar' factors
                        Gbar <- another_experiment$.Gbar
                        Qbar <- another_experiment$.Qbar
                        ## sampling
            ## ## context
            params <- formals(another_experiment$.QW)
            W <- stats::runif(n, min = eval(params$min),
                max = eval(params$max))
            ## ## action undertaken
            A <- stats::rbinom(n, size = 1, prob = Gbar(W))
            ## ## reward
            params <- formals(another_experiment$.qY)
            shape1 <- eval(params$shape1)
            QAW <- Qbar(cbind(A = A, W = W), h = h)
            Y <- stats::rbeta(n,
                                    shape1 = shape1,
                            shape2 = shape1 * (1 - QAW)/QAW)
        ## ## observation
        obs <- cbind(W = W, A = A, Y = Y)
        return(obs)
    }
#> <bytecode: Oxf3593b0>
#> <environment: Ox10837e90>
```

Let us call $\Pi_{h} \in \mathscr{M}$ the law encoded by another_experiment for a given h taken in $]-1,1[$. Note that

$$
\mathscr{P} \doteq\left\{\Pi_{h}: h \in\right]-1,1[ \}
$$

defines a collection of laws, i.e., a statistical model.
We say that $\mathscr{P}$ is a submodel of $\mathscr{M}$ because $\mathscr{P} \subset \mathscr{M}$. Moreover, we say that this submodel is through $\Pi_{0}$ since $\Pi_{h}=\Pi_{0}$ when $h=0$. We also say that $\mathscr{P}$ is a fluctuation of $\Pi_{0}$.

One could enumerate many possible submodels in $\mathscr{M}$ through $\Pi_{0}$. It turns out that all that matters for our purposes is the form of the submodel in a neighborhood of $\Pi_{0}$. We informally say that this local behavior describes the direction of a submodel through $\Pi_{0}$. We formalize this notion Section 3.3.

We now have a notion of how to move through the model space $P \in \mathscr{M}$ and can study how the value of the parameter changes as we move away from a law $P$. Above, we said that $\Psi$ is
a smooth parameter if it does not change "abruptly" as we move towards $P$ in any particular direction. That is, we should hope that $\Psi$ is differentiable along our submodel at $P$. This idea too is formalized in Section 3.3. We now turn to illustrating this idea numerically.

### 3.1.2. Numerical illustration

The code below evaluates how the parameter changes for laws in $\mathscr{P}$, and approximates the derivative of the parameter along the submodel $\mathscr{P}$ at $\Pi_{0}$. Recall that the numerical value of $\Psi\left(\Pi_{0}\right)$ has already been computed and is stored in object psi_Pi_zero.

```
approx <- seq(-1, 1, length.out = 1e2)
psi_Pi_h <- sapply(approx, function(t) {
    evaluate_psi(another_experiment, h = t)
})
slope_approx <- (psi_Pi_h - psi_Pi_zero) / approx
slope_approx <- slope_approx[min(which(approx > 0))]
ggplot() +
    geom_point(data = data.frame(x = approx, y = psi_Pi_h), aes(x, y),
                            color = "#CC6666") +
    geom_segment(aes(x = -1, y = psi_Pi_zero - slope_approx,
                    xend = 1, yend = psi_Pi_zero + slope_approx),
            arrow = arrow(length = unit(0.03, "npc")),
            color = "#9999CC") +
    geom_vline(xintercept = 0, color = "#66CC99") +
    geom_hline(yintercept = psi_Pi_zero, color = "#66CC99") +
    labs(x = "h", y = expression(Psi(Pi[h])))
```



FIGURE 2. Evolution of statistical mapping $\Psi$ along fluctuation $\left\{\Pi_{h}: h \in H\right\}$.

The red curve represents the function $h \mapsto \Psi\left(\Pi_{h}\right)$. The blue line represents the tangent to the previous curve at $h=0$, which indeed appears to be differentiable around $h=0$. In Section 3.4, we derive a closed-form expression for the slope of the blue curve.

### 3.2. Yet another experiment

1. Adapt the code from Problem 1 in Section 1.3 to visualize $w \mapsto \mathrm{E}_{\Pi_{h}}(Y \mid A=1, W=w)$, $w \mapsto \mathrm{E}_{\Pi_{h}}(Y \mid A=0, W=w)$, and $w \mapsto \mathrm{E}_{\Pi_{h}}(Y \mid A=1, W=w)-\mathrm{E}_{\Pi_{h}}(Y \mid A=0, W=w)$, for $h \in\{-1 / 2,0,1 / 2\}$.
2. Run the following chunk of code.
```
yet_another_experiment <- copy(another_experiment)
alter(yet_another_experiment,
    Qbar = function(AW, h){
        A <- AW[, "A"]
        W <- AW[, "W"]
        expit( logit( A * W + (1 - A) * W^2 ) +
            h * (2*A - 1) / ifelse(A == 1,
                                    sin((1 + W) * pi / 6),
                                    1 - sin((1 + W) * pi / 6)) *
            (Y - A * W + (1 - A) * W
    })
```

3. Justify that yet_another_fluctuation characterizes another fluctuation of $\Pi_{0}$. Comment upon the similarities and differences between $\left\{\Pi_{h}: h \in\right]-1,1[ \}$ and $\left\{\Pi_{h}^{\prime}: h \in\right.$ ] $-1,1[ \}$.
4. Repeat Problem 1 above with $\Pi_{h}^{\prime}$ substituted for $\Pi_{h}$.
5. Re-produce Figure 2 for the $\left\{\Pi_{h}^{\prime}: h \in\right]-1,1[ \}$ fluctuation. Comment on the similarities and differences between the resulting figure and Figure 2. In particular, how does the behavior of the target parameter around $h=0$ compare between laws $\Pi_{0}$ and $\Pi_{0}^{\prime}$ ?

### 3.3. 2 More on fluctuations and smoothness

### 3.3.1. Fluctuations

Let us now formally define what it means for statistical mapping $\Psi$ to be smooth at every $P \in \mathscr{M}$. Let $H$ be the interval $]-1 / M, M\left[\right.$. For every $h \in H$, we can define a law $P_{h} \in \mathscr{M}$ by setting $P_{h} \ll P^{6}$ and

$$
\begin{equation*}
\frac{d P_{h}}{d P} \doteq 1+h s \tag{9}
\end{equation*}
$$

where $s: \mathscr{O} \rightarrow \mathbb{R}$ is a (measurable) function of $O$ such that $s(O)$ is not equal to zero $P$-almost surely, $\mathrm{E}_{P}(s(O))=0$, and $s$ bounded by $M$. We make the observation that

$$
\begin{equation*}
\text { (i) }\left.\quad P_{h}\right|_{h=0}=P,\left.\quad(i i) \quad \frac{d}{d h} \log \frac{d P_{h}}{d P}(O)\right|_{h=0}=s(O) \tag{10}
\end{equation*}
$$

Because of $(i),\left\{P_{h}: h \in H\right\}$ is a submodel through $P$, also referred to as a fluctuation of $P$. The fluctuation is a one-dimensional submodel of $\mathscr{M}$ with univariate parameter $h \in H$. We note

[^4]that (ii) indicates that the score of this submodel at $h=0$ is $s$. Thus, we say that the fluctuation is in the direction of $s$.

Fluctuations of $P$ do not necessarily take the same form as in (9). No matter how the fluctuation is built, for our purposes the most important feature of the fluctuation is its direction.

### 3.3.2. Smoothness and gradients

We are now prepared to provide a formal definition of smoothness of statistical mappings. We say that a statistical mapping $\Psi$ is smooth at every $P \in \mathscr{M}$ if for each $P \in \mathscr{M}$, there exists a (measurable) function $D^{*}(P): \mathscr{O} \rightarrow \mathbb{R}$ such that $\mathrm{E}_{P}\left(D^{*}(P)(O)\right)=0, \operatorname{Var}_{P}\left(D^{*}(P)(O)\right)<\infty$, and, for every fluctuation $\left\{P_{h}: h \in H\right\}$ with score $s$ at $h=0$, the real-valued mapping $h \mapsto \Psi\left(P_{h}\right)$ is differentiable at $h=0$, with a derivative equal to

$$
\begin{equation*}
\mathrm{E}_{P}\left(D^{*}(P)(O) s(O)\right) \tag{11}
\end{equation*}
$$

The object $D^{*}(P)$ in (11) is called a gradient of $\Psi$ at $P .{ }^{7}$

### 3.3.3. A Euclidean perspective

This terminology has a direct parallel to directional derivatives in the calculus of Euclidean geometry. Recall that if $f$ is a differentiable mapping from $\mathbb{R}^{p}$ to $\mathbb{R}$, then the directional derivative of $f$ at a point $x$ (an element of $\mathbb{R}^{p}$ ) in direction $u$ (a unit vector in $\mathbb{R}^{p}$ ) is the scalar product of the gradient of $f$ and $u$. In words, the directional derivative of $f$ at $x$ can be represented as a scalar product of the direction that we approach $x$ and the change of the function's value at $x$.

In the present problem, the law $P$ is the point at which we evaluate the function $\Psi$, the score $s$ of the fluctuation is the direction in which we approach the point, and the gradient describes the change in the function's value at the point.

### 3.3.4. The canonical gradient

In general, it is possible for many gradients to exist ${ }^{8}$. Yet, in the special case that the model is nonparametric, only a single gradient exists. The unique gradient is then referred to as the canonical gradient or, for reasons that will be clarified in Section 3.5, the efficient influence curve. In the more general setting, the canonical gradient may be defined as the minimizer of $D \mapsto \operatorname{Var}_{P}(D(O))$ over the set of all gradients of $\Psi$ at $P$.

It is not difficult to check that the efficient influence curve of statistical mapping $\Psi$ (6) at $P \in \mathscr{M}$ can be written as

$$
\begin{equation*}
D^{*}(P) \doteq D_{1}^{*}(P)+D_{2}^{*}(P), \quad \text { where } \tag{12}
\end{equation*}
$$

[^5]\[

$$
\begin{aligned}
& D_{1}^{*}(P)(O) \doteq \bar{Q}(1, W)-\bar{Q}(0, W)-\Psi(P) \\
& D_{2}^{*}(P)(O) \doteq \frac{2 A-1}{\ell \bar{G}(A, W)}(Y-\bar{Q}(A, W))
\end{aligned}
$$
\]

A method from package tlrider evaluates the efficient influence curve at a law described by an object of class LAW. It is called evaluate_eic. For instance, the next chunk of code evaluates the efficient influence curve $D^{*}\left(P_{0}\right)$ of $\Psi(6)$ at $P_{0} \in \mathscr{M}$ that is characterized by experiment:

```
eic_experiment <- evaluate_eic(experiment)
```

The efficient influence curve $D^{*}\left(P_{0}\right)$ is a function from $\mathscr{O}$ to $\mathbb{R}$. As such, it can be evaluated at the five independent observations drawn from $P_{0}$ in Section 1.2.2. This is what the next chunk of code does:

```
(eic_experiment(five_obs))
#> [1] 0.246 -0.157-0.534 -0.243 0.234
```

Finally, the efficient influence curve can be visualized as two images that represent $(w, y) \mapsto$ $D^{*}\left(P_{0}\right)(w, a, y)$ for $a=0,1$, respectively:

```
crossing(w = seq(0, 1, length.out = 2e2),
        a = c(0, 1),
        y = seq(0, 1, length.out = 2e2)) %>%
    mutate(eic = eic_experiment(cbind(Y=y,A=a,W=w))) %>%
    ggplot(aes(x = w, y = y, fill = eic)) +
    geom_raster(interpolate = TRUE) +
    geom_contour(aes(z = eic), color = "white") +
    facet_wrap(~ a, nrow = 1,
                                labeller = as_labeller(c('0' = "a = 0", '1' = "a = 1"))) +
    labs(fill = expression(paste(D^"*", (P[0])(w,a,y))))
```


### 3.4. A fresh look at another_experiment

We can give a fresh look at Section 3.1.2 now.

### 3.4.1. Deriving the efficient influence curve

It is not difficult (though cumbersome) to verify that, up to a constant, $\left\{\Pi_{h}: h \in[-1,1]\right\}$ is a fluctuation of $\Pi_{0}$ in the direction (in the sense of (9)) of

$$
\begin{align*}
\sigma_{0}(O) \doteq & -10 \sqrt{W} A \times \beta_{0}(A, W) \\
& \times\left(\log (1-Y)+\sum_{k=0}^{3}\left(k+\beta_{0}(A, W)\right)^{-1}\right)+\text { constant, where }  \tag{13}\\
\beta_{0}(A, W) \doteq & \frac{1-\bar{Q}_{\Pi_{0}}(A, W)}{\bar{Q}_{\Pi_{0}}(A, W)} . \tag{14}
\end{align*}
$$



Figure 3. Visualizing the efficient influence curve $D^{*}\left(P_{0}\right)$ of $\Psi(6)$ at $P_{0}$, the law described by experiment.

Consequently, the slope of line in Figure 2 is equal to

$$
\begin{equation*}
\mathrm{E}_{\Pi_{0}}\left(D^{*}\left(\Pi_{0}\right)(O) \sigma_{0}(O)\right) \tag{15}
\end{equation*}
$$

Since $D^{*}\left(\Pi_{0}\right)$ is centered under $\Pi_{0}$, knowing $\sigma_{0}$ up to a constant is not problematic.

### 3.4.2. Numerical validation

In the following code, we check the above fact numerically. When we ran example(tlrider), we created a function sigma0. The function implements $\sigma_{0}$ defined in (14):

```
sigma0
#> function(obs, law = another_experiment) {
#> ## preliminary
#> Qbar <- get_feature(law, "Qbar", h = 0)
#> QAW <- Qbar(obs[, c("A", "W")])
#> params <- formals(get_feature(law, "qY", h = 0))
#> shape1 <- eval(params$shape1)
#> ## computations
#> betaAW <- shape1 * (1 - QAW) / QAW
#> out <- log(1 - obs[, "Y"])
#> for (int in 1:shape1) {
    out <- out + 1/(int - 1 + betaAW)
    }
    out <- - out * shape1 * (1 - QAW) / QAW *
    10 * sqrt(obs[, "W"]) * obs[, "A"]
    ## no need to center given how we will use it
    return(out)
}
```

The next chunk of code approximates (15) pointwise and with a confidence interval of asymptotic level 95\%:

```
eic_another_experiment <- evaluate_eic(another_experiment, h = 0)
obs_another_experiment <- sample_from(another_experiment, B, h = 0)
vars <- eic_another_experiment(obs_another_experiment) *
    sigma0(obs_another_experiment)
sd_hat <- sd(vars)
(slope_hat <- mean(vars))
#> [1] 1.36
(slope_CI <- slope_hat + c(-1, 1) *
    qnorm(1 - alpha / 2) * sd_hat / sqrt(B))
#> [1] 1.35 1.36
```

Equal to 1.349 (rounded to three decimal places - hereafter, all rounding will be to three decimal places as well), the first numerical approximation slope_approx is not too off!

### 3.5. 2 Asymptotic linearity and statistical efficiency

### 3.5.1. Asymptotic linearity

Suppose that $O_{1}, \ldots, O_{n}$ are drawn independently from $P \in \mathscr{M}$. If an estimator $\psi_{n}$ of $\Psi(P)$ can be written as

$$
\begin{equation*}
\psi_{n}=\Psi(P)+\frac{1}{n} \sum_{i=1}^{n} \operatorname{IC}\left(O_{i}\right)+o_{P}(1 / \sqrt{n}) \tag{16}
\end{equation*}
$$

for some function IC : $\mathscr{O} \rightarrow \mathbb{R}$ such that $\mathrm{E}_{P}(\mathrm{IC}(O))=0$ and $\operatorname{Var}_{P}(\mathrm{IC}(O))<\infty$, then we say that $\psi_{n}$ is asymptotically linear with influence curve IC. Asymptotically linear estimators are weakly convergent. Specifically, if $\psi_{n}$ is asymptotically linear with influence curve IC, then

$$
\begin{equation*}
\sqrt{n}\left(\psi_{n}-\Psi(P)\right)=\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \operatorname{IC}\left(O_{i}\right)+o_{P}(1) \tag{17}
\end{equation*}
$$

and, by the central limit theorem (recall that $O_{1}, \ldots, O_{n}$ are independent), $\sqrt{n}\left(\psi_{n}-\Psi(P)\right)$ converges in law to a centered Gaussian distribution with variance $\operatorname{Var}_{P}(\operatorname{IC}(O))$.

### 3.5.2. Influence curves and gradients

As it happens, influence curves of regular ${ }^{9}$ estimators are intimately related to gradients. In fact, if $\psi_{n}$ is a regular, asymptotically linear estimator of $\Psi(P)$ with influence curve IC, then it must be true that $\Psi$ is a smooth parameter mapping at $P$ and that IC is a gradient of $\Psi$ at $P$.

[^6]
### 3.5.3. Asymptotic efficiency

Now recall that, in Section 3.3.4, we defined the canonical gradient as the minimizer of $D \mapsto$ $\operatorname{Var}_{P}(D(O))$ over the set of all gradients. Therefore, if $\psi_{n}$ is a regular, asymptotically linear estimator of $\Psi(P)$ (built from $n$ independent observations drawn from $P$ ), then the asymptotic variance of $\sqrt{n}\left(\psi_{n}-\Psi(P)\right)$ cannot be smaller than the variance of the canonical gradient of $\Psi$ at $P$, i.e.,

$$
\begin{equation*}
\operatorname{Var}_{P}\left(D^{*}(P)(O)\right) . \tag{18}
\end{equation*}
$$

In other words, (18) is the lower bound on the asymptotic variance of any regular, asymptotically linear estimator of $\Psi(P)$. This bound is referred to as the Cramér-Rao bound. Any regular estimator that achieves this variance bound is said to be asymptotically efficient at $P$. Because the canonical gradient is the influence curve of an asymptotically efficient estimator, it is often referred to as the efficient influence curve.

### 3.6. Cramér-Rao bounds

1. What does the following chunk do?
```
obs <- sample_from(experiment, B)
(cramer_rao_hat <- var(eic_experiment(obs)))
#> [1] 0.287
```

2. Same question about this one.
```
obs_another_experiment <- sample_from(another_experiment, B, h = 0)
(cramer_rao_Pi_zero_hat <-
    var(eic_another_experiment(obs_another_experiment)))
#> [1] 0.0957
```

3. With a large independent sample drawn from $\Psi\left(P_{0}\right)$ (or $\Psi\left(\Pi_{0}\right)$ ), is it possible to construct a regular estimator $\psi_{n}$ of $\Psi\left(P_{0}\right)$ (or $\Psi\left(\Pi_{0}\right)$ ) such that the asymptotic variance of $\sqrt{n}$ times $\psi_{n}$ minus its target be smaller than the Cramér-Rao bound?
4. Is it easier to estimate $\Psi\left(P_{0}\right)$ or $\Psi\left(\Pi_{0}\right)$ (from independent observations drawn from either law)? In what sense? (Hint: you may want to compute a ratio.)

## 4. Double-robustness

### 4.1. Linear approximations of parameters

### 4.1.1. From gradients to estimators

We learned in Section 3 that the stochastic behavior of a regular, asymptotically linear estimator of $\Psi(P)$ can be characterized by its influence curve. Moreover, we said that this influence curve must in fact be a gradient of $\Psi$ at $P$.

In this section, we show that the converse is also true: given a gradient $D^{*}$ of $\Psi$ at $P$, under so-called regularity conditions, it is possible to construct an estimator with influence curve
equal to $D^{*}(P)$. This fact will suggest concrete strategies for generating efficient estimators of smooth parameters. We take here the first step towards generating such estimators: linearizing the parameter.

### 4.1.2. A Euclidean perspective

As in Section 3.3.3, drawing a parallel to Euclidean geometry is helpful. We recall that if $f$ is a differentiable mapping from $\mathbb{R}^{p}$ to $\mathbb{R}$, then a Taylor series approximates $f$ at a point $x_{0} \in \mathbb{R}^{p}$ :

$$
f\left(x_{0}\right) \approx f(x)+\left\langle\left(x_{0}-x\right), \nabla f(x)\right\rangle
$$

where $x$ is a point in $\mathbb{R}^{p}, \nabla f(x)$ is the gradient of $f$ evaluated at $x$ and $\langle u, v\rangle$ is the scalar product of $u, v \in \mathbb{R}^{p}$. As the squared distance $\left\|x-x_{0}\right\|^{2}=\left\langle x-x_{0}, x-x_{0}\right\rangle$ between $x$ and $x_{0}$ decreases, the linear approximation to $f\left(x_{0}\right)$ becomes more accurate.

### 4.1.3. The remainder term

Returning to the present problem with this in mind, we find that indeed a similar approximation strategy may be applied.

For clarity, let us introduce a new shorthand notation. For any measurable function $f$ of the observed data $O$, we may write from now on $P f \doteq \mathrm{E}_{P}(f(O))$. One may argue that the notation is valuable beyond the gain of space. For instance, (17)

$$
\sqrt{n}\left(\psi_{n}-\Psi(P)\right)=\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \operatorname{IC}\left(O_{i}\right)+o_{P}(1)
$$

can be rewritten as

$$
\sqrt{n}\left(\psi_{n}-\Psi(P)\right)=\sqrt{n}\left(P_{n}-P\right) \mathrm{IC}+o_{P}(1)
$$

thus suggesting more clearly the importance of the so-called empirical process $\sqrt{n}\left(P_{n}-P\right)$.
In particular, if $\Psi$ is smooth uniformly over directions, then for any given $P \in \mathscr{M}$, we can write

$$
\begin{equation*}
\Psi\left(P_{0}\right)=\Psi(P)+\left(P_{0}-P\right) D^{*}(P)-\operatorname{Rem}_{P_{0}}(P) \tag{19}
\end{equation*}
$$

where $\operatorname{Rem}_{P_{0}}(P)$ (defined implicitly by (19) - see (20)) is a remainder term satisfying that

$$
\frac{\operatorname{Rem}_{P_{0}}(P)}{d\left(P, P_{0}\right)} \rightarrow 0 \text { as } d\left(P, P_{0}\right) \rightarrow 0
$$

with $d$ a measure of discrepancy for distributions in $\mathscr{M}$. Note that (19) can be equivalently written as

$$
\Psi\left(P_{0}\right)=\Psi(P)+\mathrm{E}_{P_{0}}\left(D^{*}(P)(O)\right)-\mathrm{E}_{P}\left(D^{*}(P)(O)\right)-\operatorname{Rem}_{P_{0}}(P)
$$

The remainder term formalizes the notion that if $P$ is close to $P_{0}$ (i.e., if $d\left(P, P_{0}\right)$ is small), then the linear approximation of $\Psi\left(P_{0}\right)$ is more accurate. In light of the Euclidean perspective of Section 4.1.2, the remainder term $\operatorname{Rem}_{P_{0}}(P)$ plays the role of the squared distance $\left\|x-x_{0}\right\|^{2}$.

### 4.1.4. Expressing the remainder term as a function of the relevant features

The equations for the definition of the parameter (6), form of the canonical gradient (12), and linearization of parameter (19) combine to determine the remainder:

$$
\begin{equation*}
\operatorname{Rem}_{P_{0}}(P) \doteq \Psi(P)-\Psi\left(P_{0}\right)-\left(P_{0}-P\right) D^{*}(P) \tag{20}
\end{equation*}
$$

hence

$$
\begin{align*}
\operatorname{Rem}_{P_{0}}(P)=\mathrm{E}_{P_{0}}\left[\left(\bar{G}_{0}(W)-\bar{G}(W)\right)\right. & \\
& \left.\times\left(\frac{\bar{Q}_{0}(1, W)-\bar{Q}(1, W)}{\ell \bar{G}(1, W)}+\frac{\bar{Q}_{0}(0, W)-\bar{Q}(0, W)}{\ell \bar{G}(0, W)}\right)\right] \tag{21}
\end{align*}
$$

Acting as oracles, we can compute explicitly the remainder term $\operatorname{Rem}_{P_{0}}(P)$. The evaluate_remainder method makes it very easy (simply run ?evaluate_remainder to see the man page of the method):

```
(evaluate_remainder(experiment, experiment))
#> [1] 0
(rem <- evaluate_remainder(experiment, another_experiment,
    list(list(), list(h = 0))))
```

```
#> [1] 0.199
```

```
#> [1] 0.199
```

We recover the equality $\operatorname{Rem}_{P_{0}}\left(P_{0}\right)=0$, which is fairly obvious given (19). In addition, we learn that $\operatorname{Rem}_{P_{0}}\left(\Pi_{0}\right)$ equals 0.199 . In the next subsection, we invite you to make better acquaintance with the remainder term by playing around with it numerically.

### 4.2. The remainder term

1. Compute numerically $\operatorname{Rem}_{\Pi_{0}}\left(\Pi_{h}\right)$ for $h \in[-1,1]$ and plot your results. What do you notice?
2. 2 Approximate $\operatorname{Rem}_{P_{0}}\left(\Pi_{0}\right)$ numerically without relying on method evaluate_remainder and compare the value you get with that of rem. (Hint: use (20) and a large sample of observations drawn independently from $P_{0}$.)

### 4.3. 2 Double-robustness

### 4.3.1. The key property

Let us denote by $\|f\|_{P}^{2}$ the square of the $L^{2}(P)$-norm of any function $f$ from $\mathscr{O}$ to $\mathbb{R}$ i.e., using a recently introduced notation, $\|f\|_{P}^{2} \doteq P f^{2}$. For instance, $\left\|\bar{Q}_{1}-\bar{Q}_{0}\right\|_{P}$ or $\left\|\bar{G}_{1}-\bar{G}_{0}\right\|_{P}$ is a distance separating the features $\bar{Q}_{1}$ and $\bar{Q}_{0}$ or $\bar{G}_{1}$ and $\bar{G}_{0}$.

The efficient influence curve $D^{*}(P)$ at $P \in \mathscr{M}$ enjoys a rather remarkable property: it is doublerobust. Specifically, for every $P \in \mathscr{M}$, the remainder term $\operatorname{Rem}_{P_{0}}(P)$ satisfies

$$
\begin{equation*}
\operatorname{Rem}_{P_{0}}(P)^{2} \leq\left\|\bar{Q}-\bar{Q}_{0}\right\|_{P_{0}}^{2} \times\left\|\left(\bar{G}-\bar{G}_{0}\right) / \ell \bar{G}_{0}\right\|_{P_{0}}^{2} \tag{22}
\end{equation*}
$$

where $\bar{Q}$ and $\bar{G}$ are the counterparts under $P$ to $\bar{Q}_{0}$ and $\bar{G}_{0}$. The proof consists in a straightforward application of the Cauchy-Schwarz inequality to the right-hand side expression in (20).

### 4.3.2. Its direct consequence

It may not be clear yet why (22) is an important property, and why $D^{*}$ is said double-robust because of it. To answer the latter question, let us consider a law $P \in \mathscr{M}$ such that either $\bar{Q}=\bar{Q}_{0}$ or $\bar{G}=\bar{G}_{0}$.

It is then the case that either $\left\|\bar{Q}-\bar{Q}_{0}\right\|_{P}=0$ or $\left\|\bar{G}-\bar{G}_{0}\right\|_{P}=0$. Therefore, in light of (22), it also holds that $\operatorname{Rem}_{P_{0}}(P)=0 .{ }^{10}$ It thus appears that (19) simplifies to

$$
\begin{aligned}
\Psi\left(P_{0}\right) & =\Psi(P)+\left(P_{0}-P\right) D^{*}(P) \\
& =\Psi(P)+P_{0} D^{*}(P)
\end{aligned}
$$

where the second equality holds because $P D^{*}(P)=0$ for all $P \in \mathscr{M}$ by definition of $D^{*}(P)$.
It is now clear that for such a law $P \in \mathscr{M}, \Psi(P)=\Psi\left(P_{0}\right)$ is equivalent to

$$
\begin{equation*}
P_{0} D^{*}(P)=0 \tag{23}
\end{equation*}
$$

Most importantly, in words, if $P$ solves the so-called $P_{0}$-specific efficient influence curve equation (23) and if, in addition, $P$ has the same $\bar{Q}$-feature or $\bar{G}$-feature as $P_{0}$, then $\Psi(P)=\Psi\left(P_{0}\right)$.

The conclusion is valid no matter how $P$ may differ from $P_{0}$ otherwise, hence the notion of being double-robust. This property is useful to build consistent estimators of $\Psi(P)$, as we shall see in Section 5.

### 4.4. Double-robustness

1. Go back to Problem 1 in 4.2. In light of Section 4.3, what is happening?
2. Create a copy of experiment and replace its Gbar feature with some other function of $W$ (see ?copy, ?alter and Problem 2 in Section 3.2). Call $P^{\prime}$ the element of model $\mathscr{M}$ thus characterized. Can you guess the values of $\operatorname{Rem}_{P_{0}}\left(P^{\prime}\right), \Psi\left(P^{\prime}\right)$ and $P_{0} D^{*}\left(P^{\prime}\right)$ ? Support your argument.

## 5. Inference

### 5.1. Where we stand

In the previous sections, we analyzed our target parameter and presented relevant theory for understanding the statistical properties of certain types of estimators of the parameter. The theory is also relevant for building and comparing a variety of estimators.

We assume from now on that we have available a sample $O_{1}, \ldots, O_{B}$ of independent observations drawn from $P_{0}$. This is literally the case!, and the observations are stored in obs that we created in Section 3.6.

[^7]```
iter <- 1e3
```

Equal to 1 million, the sample size B is very large. We will in fact use 1000 disjoint subsamples composed of $n$ independent observations among $O_{1}, \ldots, O_{B}$, where $n$ equals B/iter, i.e., 1000 . We will thus be in a position to investigate the statistical properties of every estimation procedure by replicating it independently 1000 times.

### 5.2. Where we are going

The following sections explore different statistical paths to inferring $\psi_{0}$ or, rather (though equivalently), $\Psi\left(P_{0}\right)$.

- Section 6 presents a simple inference strategy. It can be carried out in situations where $\bar{G}_{0}$ is already known to the statistician.
- Section 7 discusses the estimation of some infinite-dimensional features of $P_{0}$. The resulting estimators are later used to estimate $\psi_{0}$.
- Section 8 extends the inference strategy discussed in Section 6 to the case that $\bar{G}_{0}$ is not known to the statistician but estimated by her. It also presents another inference strategy that relies upon the estimation of $\bar{Q}_{0}$. A theoretical analysis reveals that both strategies, called the inverse probability of treatment weighted and $G$-computation estimation methodologies, suffer from an inherent flaw.
- Section 9 builds upon the aforementioned analysis and develops a methodological workaround to circumvent the problem revealed by the analysis. It appears indeed that the flawed estimators can be corrected. However, the so-called one-step correction comes at a price that may be high in small samples.
- Section 10 also builds on the aforementioned analysis but draws a radically different conclusion from it. Instead of trying to circumvent the problem by correcting the flawed estimators of $\psi_{0}$, it does so by correcting the estimators of the infinite-dimensional features of $P_{0}$ combined to estimate $\psi_{0}$. The section thus presents an instantiation of the general targeted minimum loss estimation procedure tailored to the estimation of $\psi_{0}$. It is the main destination of this ride in targeted learning territory, far from its outposts yet well into this exciting territory.


## 6. A simple inference strategy

### 6.1. A cautionary detour

Let us introduce first the following estimator:

$$
\begin{align*}
\psi_{n}^{a} & \doteq \frac{\mathrm{E}_{P_{n}}(A Y)}{\mathrm{E}_{P_{n}}(A)}-\frac{\mathrm{E}_{P_{n}}((1-A) Y)}{\mathrm{E}_{P_{n}}(1-A)} \\
& =\frac{\sum_{i=1}^{n} \mathbf{1}\left\{A_{i}=Y_{i}=1\right\}}{\sum_{i=1}^{n} \mathbf{1}\left\{A_{i}=1\right\}}-\frac{\sum_{i=1}^{n} \mathbf{1}\left\{A_{i}=0, Y_{i}=1\right\}}{\sum_{i=1}^{n} \mathbf{1}\left\{A_{i}=0\right\}} \tag{24}
\end{align*}
$$

Note that $\psi_{n}^{a}$ is simply the difference in sample means between observations with $A=1$ and observations with $A=0$. It estimates

$$
\begin{aligned}
\Phi\left(P_{0}\right) & \doteq \frac{\mathrm{E}_{P_{0}}(A Y)}{\mathrm{E}_{P_{0}}(A)}-\frac{\mathrm{E}_{P_{0}}((1-A) Y)}{\mathrm{E}_{P_{0}}(1-A)} \\
& =\mathrm{E}_{P_{0}}(Y \mid A=1)-\mathrm{E}_{P_{0}}(Y \mid A=0)
\end{aligned}
$$

We seize this opportunity to demonstrate numerically that $\psi_{n}^{a}$ does not estimate $\Psi\left(P_{0}\right)$ because, in general, $\Psi\left(P_{0}\right)$ and $\Phi\left(P_{0}\right)$ differ. This is apparent in the following alternative expression of $\Phi\left(P_{0}\right)$ :

$$
\begin{aligned}
\Phi\left(P_{0}\right) & \left.=\mathrm{E}_{P_{0}}\left(\mathrm{E}_{P_{0}}(Y \mid A, W) \mid A=1\right)\right)-\mathrm{E}_{P_{0}}\left(\mathrm{E}_{P_{0}}(Y \mid A, W) \mid A=0\right) \\
& =\int \bar{Q}_{0}(1, w) d P_{0, W \mid A=1}(w)-\int \bar{Q}_{0}(0, w) d P_{0, W \mid A=0}(w)
\end{aligned}
$$

Contrast the above equalities and (1). In the latter, the outer integral is against the marginal law of $W$ under $P_{0}$. In the former, the outer integrals are respectively against the conditional laws of $W$ given $A=1$ and $A=0$ under $P_{0}$. Thus $\Psi\left(P_{0}\right)$ and $\Phi\left(P_{0}\right)$ will enjoy equality when the conditional laws of $W$ given $A=1$ and $W$ given $A=0$ both equal the marginal law of $W$. This can sometimes be ensured by design, as in an experiment where $A$ is randomly allocated to observations, irrespective of $W$. However, barring this level of experimental control, in general the naive estimator may lead us astray in our goal of drawing causal inference.

### 6.2. Delta-method

Consider the next chunk of code:

```
compute_irrelevant_estimator <- function(obs) {
    obs <- as_tibble(obs)
    Y <- pull(obs, Y)
    A <- pull(obs, A)
    psi_n <- mean(A * Y) / mean(A) - mean((1 - A) * Y) / mean(1 - A)
    Var_n <- cov(cbind(A * Y, A, (1 - A) * Y, (1 - A)))
    phi_n <- c(1 / mean(A), -mean(A * Y) / mean(A) 2 2,
            -1 / mean(1 - A),
            mean((1 - A) * Y) / mean(1 - A)^2)
    var_n <- as.numeric(t(phi_n) %*% Var_n %*% phi_n)
    sig_n <- sqrt(var_n / nrow(obs))
    tibble(psi_n = psi_n, sig_n = sig_n)
}
```

Function compute_irrelevant_estimator computes the estimator $\psi_{n}^{a}(24)$ based on the data set in obs.

Set $X_{n} \doteq n^{-1} \sum_{i=1}^{n}\left(A_{i} Y_{i}, A_{i},\left(1-A_{i}\right) Y_{i}, 1-A_{i}\right)^{\top}$ and $X \doteq(A Y, A,(1-A) Y, 1-A)^{\top}$. It happens that $X_{n}$ is asymptotically Gaussian: as $n$ goes to infinity,

$$
\sqrt{n}\left(X_{n}-\mathrm{E}_{P_{0}}(X)\right)
$$

converges in law to the centered Gaussian law with covariance matrix

$$
V_{0} \doteq \mathrm{E}_{P_{0}}\left(\left(X-\mathrm{E}_{P_{0}}(X)\right) \times\left(X-\mathrm{E}_{P_{0}}(X)\right)^{\top}\right) .
$$

Let $f: \mathbb{R} \times \mathbb{R}^{*} \times \mathbb{R} \times \mathbb{R}^{*}$ be given by $f(r, s, t, u)=r / s-t / u$. The function is differentiable.

1. Check that $\psi_{n}^{a}=f\left(X_{n}\right)$. Point out the line where $\psi_{n}^{a}$ is computed in the body of compute_irrelevant_estimator. Also point out to the line where the above asymptotic variance of $X_{n}$ is estimated with its empirical counterpart, say $V_{n}$.
2. 2 Argue how the delta-method yields that $\sqrt{n}\left(\psi_{n}^{a}-\Phi\left(P_{0}\right)\right)$ converges in law to the centered Gaussian law with a variance that can be estimated with

$$
\begin{equation*}
v_{n}^{a} \doteq \nabla f\left(X_{n}\right) \times V_{n} \times \nabla f\left(X_{n}\right)^{\top} . \tag{25}
\end{equation*}
$$

3. Check that the gradient $\nabla f$ of $f$ is given by $\nabla f(r, s, t, u) \doteq\left(1 / s,-r / s^{2},-1 / u, t / u^{2}\right)$. Point out to the line where the asymptotic variance of $\psi_{n}^{a}$ is estimated.

For instance,

```
(compute_irrelevant_estimator(head(obs, 1e3)))
#> # A tibble: 1 x 2
#> psi_n sig_n
#> <dbl> <dbl>
#> 1 0.110 0.0173
```

computes the numerical values of $\psi_{n}^{a}$ and $v_{n}^{a}$ based on the 1000 first observations in obs.

### 6.3. IPTW estimator assuming the mechanism of action known

### 6.3.1. A simple estimator

Let us assume for a moment that we know $\bar{G}_{0}$. This would have been the case indeed if we had experimental control over the actions taken by observational units, as in a randomized experiment. Note that, on the contrary, assuming $\bar{Q}_{0}$ known would be difficult to justify.

```
Gbar <- get_feature(experiment, "Gbar")
```

The alternative expression (7) suggests to estimate $\Psi\left(P_{0}\right)$ with

$$
\begin{align*}
\psi_{n}^{b} & \doteq \mathrm{E}_{P_{n}}\left(\frac{2 A-1}{\ell \bar{G}_{0}(A, W)} Y\right)  \tag{26}\\
& =\frac{1}{n} \sum_{i=1}^{n}\left(\frac{2 A_{i}-1}{\ell \bar{G}_{0}\left(A_{i}, W_{i}\right)} Y_{i}\right) . \tag{27}
\end{align*}
$$

Note how $P_{n}$ is substituted for $P_{0}$ in (27) relative to (7). However, we cannot call $\psi_{n}^{b}$ a substitution estimator because it does not write as $\Psi$ evaluated at an element of $\mathscr{M}$. This said, it is
dubbed an IPTW (inverse probability of treatment weighted) estimator because of the denominators $\ell \bar{G}_{0}\left(A_{i}, W_{i}\right)$ in its definition. ${ }^{11}$

In Section 8.2, we develop another IPTW estimator that does not assume that $\bar{G}_{0}$ is known beforehand.

### 6.3.2. Elementary statistical properties

It is easy to check that $\psi_{n}^{b}$ estimates $\Psi\left(P_{0}\right)$ consistently, but this is too little to request from an estimator of $\psi_{0}$. Better, $\psi_{n}^{b}$ also satisfies a central limit theorem: $\sqrt{n}\left(\psi_{n}^{b}-\psi_{0}\right)$ converges in law to a centered Gaussian law with asymptotic variance

$$
v^{b} \doteq \operatorname{Var}_{P_{0}}\left(\frac{2 A-1}{\ell \bar{G}_{0}(A, W)} Y\right),
$$

where $v^{b}$ can be consistently estimated by its empirical counterpart

$$
\begin{align*}
v_{n}^{b} & \doteq \operatorname{Var}_{P_{n}}\left(\frac{2 A-1}{\ell \bar{G}_{0}(A, W)} Y\right)  \tag{28}\\
& =\frac{1}{n} \sum_{i=1}^{n}\left(\frac{2 A_{i}-1}{\ell \bar{G}_{0}\left(A_{i}, W_{i}\right)} Y_{i}-\psi_{n}^{b}\right)^{2} . \tag{29}
\end{align*}
$$

We investigate empirically the statistical behavior of $\psi_{n}^{b}$ in Section 6.3.3.

### 6.3.3. Empirical investigation

The package tlrider includes compute_iptw, a function that implements the IPTW inference strategy (simply run ?compute_iptw to see its man page). For instance,

```
(compute_iptw(head(obs, 1e3), Gbar))
#> # A tibble: 1 x 2
#> psi_n sig_n
#> <dbl> <dbl>
#> 1 0.0713 0.0556
```

computes the numerical values of $\psi_{n}^{b}$ and $v_{n}^{b}$ based upon the 1000 first observations contained in obs.

The next chunk of code investigates the empirical behaviors of estimators $\psi_{n}^{a}$ and $\psi_{n}^{b}$. As explained in Section 5, we first make iter data sets out of the obs data set (second line), then build the estimators on each of them (fourth and fifth lines). After the first series of commands the object psi_hat_ab, a tibble, contains 2000 rows and four columns. For each smaller data set (identified by its id), two rows contain the values of either $\psi_{n}^{a}$ and $\sqrt{v_{n}^{a}} / \sqrt{n}$ (if type equals a) or $\psi_{n}^{b}$ and $\sqrt{v_{n}^{b}} / \sqrt{n}$ (if type equals b).

After the second series of commands, the object psi_hat_ab contains, in addition, the values of the recentered (with respect to $\psi_{0}$ ) and renormalized $\sqrt{n} / \sqrt{v_{n}^{a}}\left(\psi_{n}^{a}-\psi_{0}\right)$ and $\sqrt{n} / \sqrt{v_{n}^{b}}\left(\psi_{n}^{b}-\right.$

[^8]$\psi_{0}$ ), where $v_{n}^{a}(25)$ and $v_{n}^{b}(28)$ estimate the asymptotic variances of $\psi_{n}^{a}$ and $\psi_{n}^{b}$, respectively. Finally, bias_ab reports amounts of bias (at the renormalized scale).

We will refer to the recentering (always with respect to $\psi_{0}$ ) then renormalizing procedure as a renormalization scheme, because what may vary between two such procedures is the renormalization factor. We will judge whether or not the renormalization scheme is adequate for instance by comparing kernel density estimators of the law of estimators gone through a renormalization procedure with the standard normal density, as in Figure 4 below.

```
psi_hat_ab <- obs %>% as_tibble %>%
    mutate(id = (seq_len(n()) - 1) %% iter) %>%
    nest(obs = c(W, A, Y)) %>%
    mutate(est_a = map(obs, ~ compute_irrelevant_estimator(.)),
            est_b = map(obs, ~ compute_iptw(as.matrix(.), Gbar))) %>%
    pivot_longer(c('est_a', 'est_b'),
                names_to = "type", values_to = "estimates") %>%
    extract(type, "type", "_([ab])$") %>%
    unnest(estimates) %>% select(-obs)
(psi_hat_ab)
#> # A tibble: 2,000 x 4
#> id type psi_n sig_n
#> <dbl> <chr> <dbl> <dbl>
#> 1 0 a 0.0942 0.0171
#> 2 0 b 0.0169 0.0573
#> 3 1 a 0.105 0.0164
#> 4 1 b 0.0520}00.055
#> 5 2 a 
#> 6 2 b 0.135 0.0546
#> # ... with 1,994 more rows
psi_hat_ab <- psi_hat_ab %>%
    group_by(id) %>%
    mutate(clt = (psi_n - psi_zero) / sig_n)
(psi_hat_ab)
#> # A tibble: 2,000 x 5
#> # Groups: id [1,000]
#> id type psi_n sig_n clt
#> <dbl> <chr> <dbl> <dbl> <dbl>
#> 1 0 a 0.0942 0.0171 0.643
#> 2 0 b 0.0169 0.0573 -1.16
#> 3 1 1 a 
#> 4 1 b 0.0520 0.0551 -0.565
#> 5 2 a 
#> 6 2 b 0.135}00.0546 0.952
#> # ... with 1,994 more rows
(bias_ab <- psi_hat_ab %>%
    group_by(type) %>% summarise(bias = mean(clt)))
```

```
#> # A tibble: 2 x 2
#> type bias
#> <chr> <dbl>
#> 1 a 1.40
#> 2 b 0.0353
fig_bias_ab <- ggplot() +
    geom_line(aes(x = x, y = y),
            data = tibble(x = seq(-3, 3, length.out = 1e3),
                                    y = dnorm(x)),
            linetype = 1, alpha = 0.5) +
    geom_density(aes(clt, fill = type, colour = type),
                psi_hat_ab, alpha = 0.1) +
    geom_vline(aes(xintercept = bias, colour = type),
            bias_ab, size = 1.5, alpha = 0.5)
fig_bias_ab +
    labs(y = "",
        x = bquote(paste(sqrt(n/v[n]^{list(a, b)})*
                            (psi[n] {list(a, b)} - psi[0]))))
```



Figure 4. Kernel density estimators of the law of two estimators of $\psi_{0}$ (recentered with respect to $\psi_{0}$, and renormalized), one of them misconceived (a), the other assuming that $\bar{G}_{0}$ is known (b). Built based on iter independent realizations of each estimator.

By the above chunk of code, the averages of $\sqrt{n / v_{n}^{a}}\left(\psi_{n}^{a}-\psi_{0}\right)$ and $\sqrt{n / v_{n}^{b}}\left(\psi_{n}^{b}-\psi_{0}\right)$ computed across the realizations of the two estimators are respectively equal to 1.404 and 0.035 (see bias_ab). Interpreted as amounts of bias, those two quantities are represented by vertical lines in Figure 4. The red and blue bell-shaped curves represent the empirical laws of $\psi_{n}^{a}$ and $\psi_{n}^{b}$ (recentered with respect to $\psi_{0}$, and renormalized) as estimated by kernel density estimation. The latter is close to the black curve, which represents the standard normal density.

## 7. Nuisance parameters

### 7.1. Anatomy of an expression

From now, all the inference strategies that we will present unfold in two or three stages. For all of them, the first stage consists in estimating a selection of features of the law $P_{0}$ of the experiment. Specifically, the features are chosen among $Q_{0, W}$ (the marginal law of $W$ under $P_{0}$ ), $\bar{G}_{0}$ (the conditional probability that $A=1$ given $W$ under $P_{0}$ ) and $\bar{Q}_{0}$ (the conditional mean of $Y$ given $A$ and $W$ under $P_{0}$ ).

In this context, because they are not the parameter of primary interest (i.e., they are not the real-valued feature $\Psi\left(P_{0}\right)$ ), they are often referred to as nuisance parameters of $P_{0}$. The unflaterring expression conveys the notion that their estimation is merely an intermediate step along our path towards an inference of the target parameter.

As for the reason why $Q_{0, W}, \bar{G}_{0}$ and $\bar{Q}_{0}$ are singled out, it is because of their role in the definition of $\Psi$ and the efficient influence curve $D^{*}\left(P_{0}\right)$.

### 7.2. An algorithmic stance

In general, we can view an estimator of any feature $f_{0}$ of $P_{0}$ as the output of an algorithm $\widehat{\mathscr{A}}$ that maps any element of

$$
\mathscr{M}^{\mathrm{empirical}} \doteq\left\{\frac{1}{m} \sum_{i=1}^{m} \operatorname{Dirac}\left(o_{i}\right): m \geq 1, o_{1}, \ldots, o_{m} \in[0,1] \times\{0,1\} \times[0,1]\right\}
$$

to the set $\mathscr{F}$ where $f_{0}$ is known to live. Here, $\mathscr{M}^{\text {empirical }}$ can be interpreted as the set of all possible empirical measures summarizing the outcomes of any number of replications of the experiment $P_{0}$. In particular, $P_{n}$ belongs to this set.

The tlrider package includes such template algorithms for the estimation of $Q_{0, W}, \bar{G}_{0}$ and $\bar{Q}_{0}$. We illustrate how they work and their use in the next sections.

### 7.3. QW

For instance, estimate_QW is an algorithm $\widehat{\mathscr{A}}_{Q_{W}}$ for the estimation of the marginal law of $W$ under $P_{0}$ (to see its man page, simply run ?estimate_QW). It is a map from $\mathscr{M}^{\text {empirical }}$ to the set of laws on $[0,1]$. The following chunk of code estimates $Q_{0, W}$ based on the $n=1000$ first observations in obs:

```
QW_hat <- estimate_QW(head(obs, 1e3))
```

It is easy to sample independent observations from QW_hat. To do so, we create an object of class LAW then set its marginal law of $W$ to that described by QW_hat and specify its sample_from feature:

```
empirical_experiment <- LAW()
alter(empirical_experiment, QW = QW_hat)
alter(empirical_experiment, sample_from = function(n) {
```

```
    QW <- get_feature(empirical_experiment, "QW")
    W <- sample(pull(QW, "value"), n, prob = pull(QW, "weight"))
    cbind(W = W, A = NA, Y = NA)
})
W <- sample_from(empirical_experiment, 1e3) %>% as_tibble
W %>%
ggplot() +
    geom_histogram(aes(x = W, y = stat(density)), bins = 40) +
    stat_function(fun = get_feature(experiment, "QW"), col = "red")
```



Figure 5. Histogram representing 1000 observations drawn independently from QW_hat. The superimposed red curve is the true density of $Q_{0, W}$.

Note that all the $W \mathrm{~s}$ sampled from QW_hat fall in the set $\left\{W_{1}, \ldots, W_{n}\right\}$ of observed $W \mathrm{~s}$ in obs (an obvious fact given the definition of the sample_from feature of empirical_experiment:

```
(length(intersect(pull(W, W), head(obs[, "W"], 1e3))))
#> [1] 1000
```

This is because estimate_QW estimates $Q_{0, W}$ with its empirical counterpart, i.e.,

$$
\frac{1}{n} \sum_{i=1}^{n} \operatorname{Dirac}\left(W_{i}\right)
$$

### 7.4. Gbar

Another template algorithm is built-in into tlrider: estimate_Gbar (to see its man page, simply run ?estimate_Gbar). Unlike estimate_QW, estimate_Gbar needs further specification of the algorithm. The package also includes examples of such specifications.

There are two sorts of specifications, of which we say that they are either working modelbased or machine learning-based. We discuss the former sort in the next subsection. The latter sort is discussed in Section 7.6.

### 7.4.1. Working model-based algorithms

Let us take a look at working_model_G_one for instance:

```
working_model_G_one
#> $model
#> function (...)
#> {
#> trim_glm_fit(glm(family = binomial(), ...))
#> }
#> <environment: Oxecf7498>
#>
#> $formula
#> A ~ I(W`0.5) + I(abs(W - 5/12)^0.5) + I(W^1) + I(abs(W - 5/12)^1) +
#> I(W^1.5) + I(abs(W - 5/12)^1.5)
#> <environment: Oxecf7498>
#>
#> $type_of_preds
#> [1] "response"
#>
#> attr(,"ML")
#> [1] FALSE
```

and focus on its model and formula attributes. The former relies on the glm and binomial functions from base R, and on trim_glm_fit (which removes information that we do not need from the standard output of glm, simply run ?trim_glm_fit to see the function's man page). The latter is a formula that characterizes what we call a working model for $\bar{G}_{0}$.

In words, by using working_model_G_one we implicitly choose the so-called logistic (or negative binomial) loss function $L_{a}$ given by

$$
\begin{equation*}
-L_{a}(f)(A, W) \doteq A \log f(W)+(1-A) \log (1-f(W)) \tag{30}
\end{equation*}
$$

for any function $f:[0,1] \rightarrow[0,1]$ paired with the working model

$$
\mathscr{F}_{1} \doteq\left\{f_{\theta}: \theta \in \mathbb{R}^{7}\right\}
$$

where, for any $\theta \in \mathbb{R}^{7}$,

$$
\operatorname{logit} f_{\theta}(W) \doteq \theta_{0}+\sum_{j=1}^{3}\left(\theta_{j} W^{j / 2}+\theta_{3+j}|W-5 / 12|^{j / 2}\right)
$$

We acted as oracles when we specified the working model: it is well-specified, i.e., it happens that $\bar{G}_{0}$ is the unique minimizer of the risk entailed by $L_{a}$ over $\mathscr{F}_{1}$ :

$$
\bar{G}_{0}=\underset{f_{\theta} \in \mathscr{F}_{1}}{\arg \min } \mathrm{E}_{P_{0}}\left(L_{a}\left(f_{\theta}\right)(A, W)\right) .
$$

Therefore, the estimator $\bar{G}_{n}$ obtained by minimizing the empirical risk

$$
\mathrm{E}_{P_{n}}\left(L_{a}\left(f_{\theta}\right)(A, W)\right)=\frac{1}{n} \sum_{i=1}^{n} L_{a}\left(f_{\theta}\right)\left(A_{i}, W_{i}\right)
$$

over $\mathscr{F}_{1}$ estimates $\bar{G}_{0}$ consistently.
Of course, it is seldom certain in real life that the target feature, here $\bar{G}_{0}$, belongs to the working model. ${ }^{12}$ Suppose for instance that we choose a small finite-dimensional working model $\mathscr{F}_{2}$ without acting as an oracle. Then consistency certainly fails to hold. However, if $\bar{G}_{0}$ can nevertheless be projected unambiguously onto $\mathscr{F}_{2}$ (an assumption that cannot be checked), then the estimator might converge to the projection.

### 7.4.2. Visualization

To illustrate the use of the algorithm $\widehat{\mathscr{A}}_{\bar{G}, 1}$ obtained by combining estimate_Gbar and working_model_G_one, let us estimate $\bar{G}_{0}$ based on the first $n=1000$ observations in obs:

```
Gbar_hat <- estimate_Gbar(head(obs, 1e3), algorithm = working_model_G_one)
```

Using compute_Gbar_hat_W ${ }^{13}$ (simply run ?compute_Gbar_hat_W to see its man page) makes it is easy to compare visually the estimator $\bar{G}_{n} \doteq \widehat{\mathscr{A}}_{\bar{G}, 1}\left(P_{n}\right)$ with its target $\bar{G}_{0}$ :

```
tibble(w = seq(0, 1, length.out = 1e3)) %>%
    mutate("truth" = Gbar(w),
        "estimated" = compute_Gbar_hatW(w, Gbar_hat)) %>%
    pivot_longer(-w, names_to = "f", values_to = "value") %>%
    ggplot() +
    geom_line(aes(x = w, y = value, color = f), size = 1) +
    labs(y = "f(w)",
            title = bquote("Visualizing" ~ bar(G) [0] ~ "and" ~ hat(G)[n])) +
    ylim(NA, 1)
```


### 7.5. Qbar, working model-based algorithms

A third template algorithm is built-in into tlrider: estimate_Qbar (to see its man page, simply run ?estimate_Qbar). Like estimate_Gbar, estimate_Qbar needs further specification of the algorithm. The package also includes examples of such specifications, which can also be either working model-based (see Section 7.4) or machine learning-based (see Sections 7.6 and 7.7).

There are built-in specifications similar to working_model_G_one, e.g.,

[^9]

FIGURE 6. Comparing $\bar{G}_{n} \doteq \widehat{\mathscr{A}}_{\overline{\bar{G}}, 1}\left(P_{n}\right)$ and $\bar{G}_{0}$. The estimator is consistent because the algorithm relies on a working model that is correctly specified.

```
working_model_Q_one
#> $model
#> function (...)
#> {
#> trim_glm_fit(glm(family = binomial(), ...))
#> }
#> <environment: Oxecfr498>
#>
#> $formula
#> Y~A* (I(W^0.5) + I(W^1) + I(W^1.5))
#> <environment: Oxecfr498>
#>
#> $type_of_preds
#> [1] "response"
#>
#> attr(, "ML")
#> [1] FALSE
#> attr(,"stratify")
#> [1] FALSE
```

1. Drawing inspiration from Section 7.4 , comment upon and use the algorithm $\widehat{\mathscr{A}}_{\bar{Q}, 1}$ obtained by combining estimate_Gbar and working_model_Q_one.

### 7.6. Qbar

### 7.6.1. Qbar, machine learning-based algorithms

We explained how algorithm $\widehat{\mathscr{A}}_{\bar{G}, 1}$ is based on a working model (and you did for $\widehat{\mathscr{A}}_{\bar{Q}, 1}$ ). It is not the case that all algorithms are based on working models in the same (admittedly rather narrow)
sense. We propose to say that those algorithms that are not based on working models like $\widehat{\mathscr{A}}_{\bar{G}, 1}$, for instance, are instead machine learning-based.

Typically, machine learning-based algorithms are more data-adaptive; they rely on larger working models, and/or fine-tune parameters that must be calibrated, e.g. by cross-validation. Furthermore, they call for being stacked, i.e., combined by means of another outer algorithm (involving cross-validation) into a more powerful machine learning-based meta-algorithm. The super learning methodology is a popular stacking algorithm.

We will elaborate further on this important topic in another forthcoming part and merely touch upon it in Section 7.8. Here, we simply illustrate the concept with two specifications built-in into tlrider. Based on the $k$-nearest neighbors non-parametric estimating methodology, the first one is discussed in the next subsection. Based on boosted trees, another non-parametric estimating methodology, the second one is used in the exercise that follows the next subsection.

### 7.6.2. Qbar, kNN algorithm

Algorithm $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$ is obtained by combining estimate_Qbar and kknn_algo. The training of $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$ (i.e., the making of the output $\widehat{\mathscr{A}}_{\overline{\mathrm{Q}}, \mathrm{kNN}}\left(P_{n}\right)$ is implemented based on function caret: : train of the caret (classification and regression training) package (to see its man page, simply run ?caret: :train). Some additional specifications are provided in kknn_grid and kknn_control.

In a nutshell, $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$ estimates $\bar{Q}_{0}(1, \cdot)$ and $\bar{Q}_{0}(0, \cdot)$ separately. Each of them is estimated by applying the $k$-nearest neighbors methodology as it is implemented in function kknn::train.kknn from the kknn package (to see its man page, simply run ?kknn: :train.kknn). ${ }^{14}$ The following chunk of code trains algorithm $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$ on $P_{n}$ :

```
Qbar_hat_kknn <- estimate_Qbar(head(obs, 1e3),
    algorithm = kknn_algo,
    trControl = kknn_control,
    tuneGrid = kknn_grid)
```

Using compute_Qbar_hat_AW (simply run ?compute_Qbar_hat_AW to see its man page) makes it is easy to compare visually the estimator $\bar{Q}_{n, \mathrm{kNN}} \stackrel{\hat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}}{ }\left(\bar{P}_{n}\right)$ with its target $\bar{Q} 0$, see Figure 7.

```
fig <- tibble(w = seq(0, 1, length.out = 1e3),
    truth_1 = Qbar(cbind (A = 1, W = w)),
    truth_0 = Qbar(cbind(A=0, W = w)),
    kNN_1 = compute_Qbar_hatAW(1, w, Qbar_hat_kknn),
    kNN_0 = compute_Qbar_hatAW(0, w, Qbar_hat_kknn))
```

[^10]
### 7.6.3. Qbar, boosted trees algorithm

Algorithm $\widehat{\mathscr{A}}_{\bar{Q}}$, trees is obtained by combining estimate_Qbar and bstTree_algo. The training of $\widehat{\mathscr{A}}_{\bar{Q}}$,trees (i.e., the making of the output $\widehat{\mathscr{A}}_{\bar{Q}, \text { trees }}\left(P_{n}\right)$ is implemented based on function caret: :train of the caret package. Some additional specifications are provided in bstTree_grid and bstTree_control.

In a nutshell, $\widehat{\mathscr{A}}_{\bar{Q}, \text { trees }}$ estimates $\bar{Q}_{0}(1, \cdot)$ and $\bar{Q}_{0}(0, \cdot)$ separately. Each of them is estimated by boosted trees as implemented in function bst: : bst from the bst (gradient boosting) package (to see its man page, simply run ?bst : : bst). ${ }^{15}$ The following chunk of code trains algorithm $\widehat{\mathscr{A}}_{\bar{Q}, \text { trees }}$ on $P_{n}$, and reveals what are the optimal fine-tune parameters for the estimation of $\bar{Q}_{0}(1, \cdot)$ and $\bar{Q}_{0}(0, \cdot)$ :

```
Qbar_hat_trees <- estimate_Qbar(head(obs, 1e3),
                        algorithm = bstTree_algo,
                        trControl = bstTree_control,
                        tuneGrid = bstTree_grid)
Qbar_hat_trees %>% filter(a == "one") %>% pull(fit) %>%
    capture.output %>% tail(3) %>% str_wrap(width = 60) %>% cat
#> RMSE was used to select the optimal model using the smallest
#> value. The final values used for the model were mstop = 30,
#> maxdepth = 1 and nu = 0.1.
Qbar_hat_trees %>% filter(a == "zero") %>% pull(fit) %>%
    capture.output %>% tail(3) %>% str_wrap(width = 60) %>% cat
#> RMSE was used to select the optimal model using the smallest
#> value. The final values used for the model were mstop = 20,
#> maxdepth = 1 and nu = 0.2.
```

We can compare visually the estimators $\bar{Q}_{n, \mathrm{kNN}}, \bar{Q}_{n, \text { trees }} \doteq \widehat{\mathscr{A}}_{\bar{Q}, \text { trees }}\left(P_{n}\right)$ with its target $\bar{Q}_{0}$, see Figure 7. In summary, $\bar{Q}_{n, \mathrm{kNN}}$ is rather good, though very variable at the vincinity of the break points. As for $\bar{Q}_{n, \text { trees }}$, it does not seem to capture the shape of its target.

```
fig %>%
    mutate(trees_1 = compute_Qbar_hatAW(1, w, Qbar_hat_trees),
        trees_0 = compute_Qbar_hatAW(0, w, Qbar_hat_trees)) %>%
    pivot_longer(-w, names_to = "f", values_to = "value") %>%
    extract(f, c("f", "a"), "([^_]+)_([01]+)") %>%
    mutate(a = paste0("a=", a)) %>%
    ggplot +
    geom_line(aes(x = w, y = value, color = f), size = 1) +
    labs(y = "f(w)",
            title = bquote("Visualizing" ~ bar(Q)[0] ~ "and its estimators")) +
```

[^11]```
ylim(NA, 1) +
```

facet_wrap (~ a)

Visualizing $\bar{Q}_{0}$ and its estimators


Figure 7. Comparing to their target two (machine learning-based) estimators of $\bar{Q}_{0}$, one based on the $k$-nearest neighbors and the other on boosted trees.

### 7.7. 2 Qbar, machine learning-based algorithms

1. Using estimate_ $Q$, make your own machine learning-based algorithm for the estimation of $\bar{Q}_{0}$.
2. Train your algorithm on the same data set as $\widehat{\mathscr{A}}_{\overline{\mathrm{Q}}, \mathrm{kNN}}$ and $\widehat{\mathscr{A}}_{\overline{\mathrm{Q}}, \text { trees }}$. If, like $\widehat{\mathscr{A}}_{\overline{\bar{Q}}, \text { trees }}$, your algorithm includes a fine-tuning procedure, comment upon the optimal, data-driven specification.
3. Plot your estimators of $\bar{Q}_{0}(1, \cdot)$ and $\bar{Q}_{0}(0, \cdot)$ on Figure 7 .

### 7.8. Meta-learning/super learning

Without a great deal of previous experience or scientific expertise, it would have likely been difficult for us to a priori postulate which of the two above machine learning algorithms (or indeed the algorithm based on a working model) would perform better for estimation of $\bar{Q}_{0}$. Rather than committing to one single algorithm, we may instead wish to resort to meta-learning or super learning. In this approach, one specifies a library of candidate algorithms for estimating a given nuisance parameter. Cross-validation is used to determine an ensemble (e.g. convex combination) of the algorithms that yields the best fit to the underlying function. In this way, one can learn in real time which algorithms tend to fit the data best and shift attention towards those algorithms.

## 8. Two "naive" inference strategies

### 8.1. Why "naive"?

In this section, we present and discuss two strategies for the inference of $\Psi\left(P_{0}\right)$. In light of Section 7.1, both unfold in two stages. During the first stage, some features among $Q_{0, W}, \bar{G}_{0}$ and $\bar{Q}_{0}$ (the $\Psi$-specific nuisance parameters, see Section 7) are estimated. During the second stage, these estimators are used to build estimators of $\Psi\left(P_{0}\right)$.

Although the strategies sound well conceived, a theoretical analysis reveals that they lack a third stage trying to correct an inherent flaw. They are thus said naive. The analysis and a first modus operandi are presented in Section 9.1.

### 8.2. IPTW estimator

### 8.2.1. Construction and computation

In Section 6.3, we developed an IPTW estimator, $\psi_{n}^{b}$, assuming that we knew $\bar{G}_{0}$ beforehand. What if we did not? Obviously, we could estimate it and substitute the estimator of $\ell \bar{G}_{0}$ for $\ell \bar{G}_{0}$ in (27).

Let $\widehat{\mathscr{A}}_{\overline{\bar{G}}}$ be an algorithm designed for the estimation of $\bar{G}_{0}$ (see Section 7.4). We denote by $\bar{G}_{n} \doteq \widehat{\mathscr{A}}\left(P_{n}\right)$ the output of the algorithm trained on $P_{n}$, and by $\ell \bar{G}_{n}$ the resulting (empirical) function given by

$$
\ell \bar{G}_{n}(A, W) \doteq A \bar{G}_{n}(W)+(1-A)\left(1-\bar{G}_{n}(W)\right) .
$$

In light of (27), we introduce

$$
\psi_{n}^{c} \doteq \frac{1}{n} \sum_{i=1}^{n}\left(\frac{2 A_{i}-1}{\ell \bar{G}_{n}\left(A_{i}, W_{i}\right)} Y_{i}\right) .
$$

From a computational point of view, the tlrider package makes it easy to build $\psi_{n}^{c}$. Recall that

```
compute_iptw(head(obs, 1e3), Gbar)
```

implements the computation of $\psi_{n}^{b}$ based on the $n=1000$ first observations stored in obs, using the true feature $\bar{G}_{0}$ stored in Gbar, see Section 6.3.3 and the construction of psi_hat_ab. Similarly,

```
Gbar_hat <- estimate_Gbar(head(obs, 1e3), working_model_G_one)
compute_iptw(head(obs, 1e3), wrapper(Gbar_hat)) %>% pull(psi_n)
#> [1] 0.0915
```

implements (i) the estimation of $\bar{G}_{0}$ with $\bar{G}_{n} /$ Gbar_hat using algorithm $\widehat{\mathscr{A}_{\bar{G}}^{\mathbf{G}}, 1}$ (first line) then (ii) the computation of $\psi_{n}^{c}$ (second line), both based on the same observations as above.

Note how we use function wrapper (simply run ?wrapper to see its man page).

### 8.2.2. Elementary statistical properties

Because $\bar{G}_{n}$ minimizes the empirical risk over a finite-dimensional, identifiable, and wellspecified working model, $\sqrt{n}\left(\psi_{n}^{c}-\psi_{0}\right)$ converges in law to a centered Gaussian law. Moreover, the asymptotic variance of $\sqrt{n}\left(\psi_{n}^{c}-\psi_{0}\right)$ is conservatively ${ }^{16}$ estimated with

$$
\begin{aligned}
v_{n}^{c} & \doteq \operatorname{Var}_{P_{n}}\left(\frac{2 A-1}{\ell \bar{G}_{n}(A, W)} Y\right) \\
& =\frac{1}{n} \sum_{i=1}^{n}\left(\frac{2 A_{i}-1}{\ell \bar{G}_{n}\left(A_{i}, W_{i}\right)} Y_{i}-\psi_{n}^{c}\right)^{2} .
\end{aligned}
$$

We investigate empirically the statistical behavior of $\psi_{n}^{c}$ in Section 8.2.3. For an analysis of the reason why $v_{n}^{c}$ is a conservative estimator of the asymptotic variance of $\sqrt{n}\left(\psi_{n}^{c}-\psi_{0}\right)$, see there in Appendix C.1.1.

Before proceeding, let us touch upon what would have happened if we had used a less amenable algorithm $\widehat{\mathscr{A}}_{\bar{G}}$. For instance, $\widehat{\mathscr{A}}_{\bar{G}}$ could still be well-specified ${ }^{17}$ but so versatile/complex (as opposed to being based on well-behaved, finite-dimensional parametric model) that the estimator $\bar{G}_{n}$, though still consistent, would converge slowly to its target. Then, root- $n$ consistency would fail to hold. Or $\widehat{\mathscr{A}}_{\bar{G}}$ could be mis-specified and there would be no guarantee at all that the resulting estimator $\psi_{n}^{c}$ be even consistent.

### 8.2.3. Empirical investigation

Let us compute $\psi_{n}^{c}$ on the same iter $=1000$ independent samples of independent observations drawn from $P_{0}$ as in Section 6.3. As explained in Sections 5 and 6.3.3, we first make iter data sets out of the obs data set (third line), then train algorithm $\widehat{\mathscr{A}}_{\bar{G}, 1}$ on each of them (fifth to seventh lines). After the first series of commands the object learned_features_fixed_sample_size, a tibble, contains 1000 rows and three columns.

We created learned_features_fixed_sample_size to store the estimators of $\bar{G}_{0}$ for future use. We will at a later stage enrich the object, for instance by adding to it estimators of $\bar{Q}_{0}$ obtained by training different algorithms on each smaller data set.

In the second series of commands, the object psi_hat_abc is obtained by adding to psi_hat_ab (see Section 6.3.3) an 1000 by four tibble containing notably the values of $\psi_{n}^{c}$ and $\sqrt{v_{n}^{c}} / \sqrt{n}$ computed by calling compute_iptw. The object also contains the values of the recentered (with respect to $\psi_{0}$ ) and renormalized $\sqrt{n} / \sqrt{v_{n}^{c}}\left(\psi_{n}^{c}-\psi_{0}\right)$. Finally, bias_abc reports amounts of bias (at the renormalized scale).

```
learned_features_fixed_sample_size <-
    obs %>% as_tibble %>%
    mutate(id = (seq_len(n()) - 1) %% iter) %>%
    nest(obs = c(W, A, Y)) % %%
    mutate(Gbar_hat =
```

${ }^{16}$ In words, $v_{n}^{c}$ converges to an upper-bound of the true asymptotic variance.
${ }^{17}$ Well-specified e.g. in the sense that the target $\bar{G}_{0}$ of $\widehat{\mathscr{A}} \widehat{\bar{G}}$ belongs to the closure of the algorithm's image $\widehat{\mathscr{A}} \overline{\bar{G}}\left(\mathscr{M}^{\text {empirical }}\right)$ or, in other words, can be approximated arbitrarily well by an output of the algorithm.

```
        map(obs,
                            ~ estimate_Gbar(., algorithm = working_model_G_one)))
psi_hat_abc <-
    learned_features_fixed_sample_size %>%
    mutate(est_c =
                    map2(obs, Gbar_hat,
                            ~ compute_iptw(as.matrix(.x), wrapper(.y, FALSE)))) %>%
    unnest(est_c) %>% select(-Gbar_hat, -obs) %>%
    mutate(clt = (psi_n - psi_zero) / sig_n,
            type = "c") %>%
    full_join(psi_hat_ab)
(bias_abc <- psi_hat_abc %>%
    group_by(type) %>% summarise(bias = mean(clt)))
#> # A tibble: 3 x 2
#> type bias
#> <chr> <dbl>
#> 1 a 1.40
#> 2 b 0.0353
#> 3 c -0.00443
```

By the above chunk of code, the average of $\sqrt{n / v_{n}^{c}}\left(\psi_{n}^{c}-\psi_{0}\right)$ computed across the realizations is equal to -0.004 (see bias_abc). In words, the bias of $\psi_{n}^{c}$ is of the same magnitude as that of $\psi_{n}^{b}$ despite the fact that the construction of $\psi_{n}^{c}$ hinges on the estimation of $\bar{G}_{0}$ (based on the well-specified algorithm $\widehat{\mathscr{A}}_{\bar{G}, 1}$ ).

We represent the empirical laws of the recentered (with respect to $\psi_{0}$ ) and renormalized $\psi_{n}^{a}$, $\psi_{n}^{b}$ and $\psi_{n}^{c}$ in Figures 8 (kernel density estimators) and 9 (quantile-quantile plots).

```
fig_bias_ab +
    geom_density(aes(clt, fill = type, colour = type), psi_hat_abc, alpha = 0.1) +
    geom_vline(aes(xintercept = bias, colour = type),
            bias_abc, size \(=1.5\), alpha \(=0.5)+\)
    xlim(-3, 4) +
    labs(y = " ",
        \(\mathrm{x}=\operatorname{bquote}(\) paste (sqrt(n/v[n] \(\{\) list (a, b, c) \})*
                            (psi[n] \{list(a, b, c)\} - psi[0]))))
ggplot(psi_hat_abc, aes(sample = clt, fill = type, colour = type)) +
    geom_abline(intercept \(=0\), slope \(=1\), alpha \(=0.5\) ) +
    geom_qq(alpha \(=1\) )
```

Figures 8 and 9 confirm that $\psi_{n}^{c}$ behaves as well as $\psi_{n}^{b}$ in terms of bias - but remember that we acted as oracles when we chose the well-specified algorithm $\widehat{\mathscr{A}}_{\bar{G}, 1}$. They also corroborate that $v_{n}^{c}$, the estimator of the asymptotic variance of $\sqrt{n}\left(\psi_{n}^{c}-\psi_{0}\right)$, is conservative: for instance, the corresponding bell-shaped blue curve is too much concentrated around its axis of symmetry.

The actual asymptotic variance of $\sqrt{n}\left(\psi_{n}^{c}-\psi_{0}\right)$ can be estimated with the empirical variance


FIGURE 8. Kernel density estimators of the law of three estimators of $\psi_{0}$ (recentered with respect to $\psi_{0}$, and renormalized), one of them misconceived (a), one assuming that $\bar{G}_{0}$ is known $(b)$ and one that hinges on the estimation of $\bar{G}_{0}(c)$. The present figure includes Figure 4 (but the colors differ). Built based on iter independent realizations of each estimator.


Figure 9. Quantile-quantile plot of the standard normal law against the empirical laws of three estimators of $\psi_{0}$, one of them misconceived (a), one assuming that $\bar{G}_{0}$ is known (b) and one that hinges on the estimation of $\bar{G}_{0}(c)$. Built based on iter independent realizations of each estimator.
of the iter replications of the construction of $\psi_{n}^{c}$.

```
(emp_sig_n <- psi_hat_abc %>% filter(type == "c") %>%
    summarize(sd(psi_n)) %>% pull)
#> [1] 0.0175
(summ_sig_n <- psi_hat_abc %>% filter(type == "c") %>% select(sig_n) %>%
    summary)
#> sig_n
#> Min. :0.0523
#> 1st Qu.:0.0549
#> Median :0.0560
#> Mean :0.0561
#> 3rd Qu.:0.0570
#> Max. :0.0631
```

The empirical standard deviation is approximately 3.213 times smaller than the average estimated standard deviation. The estimator is conservative indeed! Furthermore, note the better fit with the density of the standard normal density of the kernel density estimator of the law of $\sqrt{n}\left(\psi_{n}^{c}-\psi_{0}\right)$ renormalized with emp_sig_n.

```
clt_c <- psi_hat_abc %>% filter(type == "c") %>%
    mutate(clt = clt * sig_n / emp_sig_n)
fig_bias_ab +
    geom_density(aes(clt, fill = type, colour = type), clt_c, alpha = 0.1) +
    geom_vline(aes(xintercept = bias, colour = type),
                    bias_abc, size = 1.5, alpha = 0.5) +
    xlim(-3, 4) +
    labs(y = "",
            x = bquote(paste(sqrt(n/v[n] {list(a, b, c)})*
                    (psi[n]^{list(a, b, c)} - psi[0]))))
```

Workaround. In a real world data-analysis, one could correct the estimation of the asymptotic variance of $\sqrt{n}\left(\psi_{n}^{c}-\psi_{0}\right)$. We could for instance derive the influence function as it is described there in Appendix C.1.1 and use the corresponding influence function-based estimator of the variance. One could also rely on the cross-validation, possibly combined with the previous workaround. Or one could also rely on the bootstrap. ${ }^{18}$ This, however, would only make sense if one knew for sure that the algorithm for the estimation of $\bar{G}_{0}$ is well-specified.

### 8.3. Investigating further the IPTW inference strategy

1. Building upon the chunks of code devoted to the repeated computation of $\psi_{n}^{b}$ and its companion quantities, construct confidence intervals for $\psi_{0}$ of (asymptotic) level $95 \%$, and check if the empirical coverage is satisfactory. Note that if the coverage was exactly $95 \%$, then the number of confidence intervals that would contain $\psi_{0}$ would follow a binomial

[^12]

Figure 10. Kernel density estimators of the law of three estimators of $\psi_{0}$ (recentered with respect to $\psi_{0}$, and renormalized), one of them misconceived ( $a$ ), one assuming that $\bar{G}_{0}$ is known (b) and one that hinges on the estimation of $\bar{G}_{0}$ and an estimator of the asymptotic variance computed across the replications (c). The present figure includes Figure 4 (but the colors differ) and it should be compared to Figure 9. Built based on iter independent realizations of each estimator.
law with parameters iter and 0.95 , and recall that function binom.test performs an exact test of a simple null hypothesis about the probability of success in a Bernoulli experiment against its three one-sided and two-sided alternatives.
2. Discuss what happens when the dimension of the (still well-specified) working model grows. Start with the built-in working model working_model_G_two. The following chunk of code re-defines working_model_G_two

```
## make sure '1/2' and '1' belong to 'powers'
powers <- rep(seq(1/4, 3, by = 1/4), each = 2)
working_model_G_two <- list(
    model = function(...) {trim_glm_fit(glm(family = binomial(), ...))},
    formula = stats::as.formula(
        paste("A ~",
            paste(c("I(W`", "I(abs(W - 5/12)^"),
                powers,
                        sep = "", collapse = ") + "),
            ")")
    ),
    type_of_preds = "response"
)
attr(working_model_G_two, "ML") <- FALSE
```

3. Play around with argument powers (making sure that $1 / 2$ and 1 belong to it), and plot graphics similar to those presented in Figures 8 and 9.
4. Discuss what happens when the working model is mis-specified. You could use the built-in working model working_model_G_three.
5. Repeat the analysis developed in response to problem 1 above but for $\psi_{n}^{c}$. What can you say about the coverage of the confidence intervals?
6. 2 (Follow-up to problem 5). Implement the bootstrap procedure evoked at the end of Section 8.2.3. Repeat the analysis developed in response to problem 1. Compare your results with those to problem 5.
7. 2 Is it legitimate to infer the asymptotic variance of $\psi_{n}^{c}$ with $v_{n}^{c}$ when one relies on a very data-adaptive/versatile algorithm to estimate $\bar{G}_{0}$ ?

### 8.4. G-computation estimator

### 8.4.1. Construction and computation

Let $\widehat{\mathscr{A}}_{Q_{W}}$ be an algorithm designed for the estimation of $Q_{0, W}$ (see Section 7.3). We denote by $Q_{n, W} \doteq \widehat{\mathscr{A}}_{Q_{W}}\left(P_{n}\right)$ the output of the algorithm trained on $P_{n}$.

Let $\widehat{\mathscr{A}}_{\bar{Q}}$ be an algorithm designed for the estimation of $\bar{Q}_{0}$ (see Section 7.6). We denote by $\bar{Q}_{n} \doteq \widehat{\mathscr{A}_{\bar{Q}}}\left(P_{n}\right)$ the output of the algorithm trained on $P_{n}$.

Equation (1) suggests the following, simple estimator of $\Psi\left(P_{0}\right)$ :

$$
\begin{equation*}
\psi_{n} \doteq \int\left(\bar{Q}_{n}(1, w)-\bar{Q}_{n}(0, w)\right) d Q_{n, W}(w) . \tag{31}
\end{equation*}
$$

In words, this estimator is implemented by first regressing $Y$ on $(A, W)$, then by marginalizing with respect to the estimated law of $W$. The resulting estimator is referred to as a $G$-computation (or standardization) estimator.

From a computational point of view, the tlrider package makes it easy to build the Gcomputation estimator. Recall that we have already estimated the marginal law $Q_{0, W}$ of $W$ under $P_{0}$ by training the algorithm $\widehat{\mathscr{A}}_{Q_{W}}$ as it is implemented in estimate_QW on the $n=1000$ first observations in obs (see Section 7.3):

QW_hat <- estimate_QW(head(obs, 1e3))
Recall that $\widehat{\mathscr{A}}_{\bar{Q}, 1}$ is the algorithm for the estimation of $\bar{Q}_{0}$ as it is implemented in estimate_Qbar with its argument algorithm set to the built-in working_model_Q_one (see Section 7.5). Recall also that $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$ is the algorithm for the estimation of $\bar{Q}_{0}$ as it is implemented in estimate_Qbar with its argument algorithm set to the built-in kknn_algo (see Section 7.6.2). We have already trained the latter on the $n=1000$ first observations in obs. Let us train the former on the same data set:

```
Qbar_hat_kknn <- estimate_Qbar(head(obs, 1e3),
    algorithm = kknn_algo,
    trControl = kknn_control,
    tuneGrid = kknn_grid)
```

Qbar_hat_d <- estimate_Qbar(head(obs, 1e3), working_model_Q_one)
With these estimators handy, computing the G-computation estimator is as simple as running the following chunk of code:

```
(compute_gcomp(QW_hat, wrapper(Qbar_hat_kknn, FALSE), 1e3))
#> # A tibble: 1 x 2
#> psi_n sig_n
#> <dbl> <dbl>
#> 1 0.0887 0.00245
(compute_gcomp(QW_hat, wrapper(Qbar_hat_d, FALSE), 1e3))
#> # A tibble: 1 x 2
#> psi_n sig_n
#> <dbl> <dbl>
#> 1 0.0961 0.00161
```

Note how we use function wrapper again, and that it is necessary to provide the number of observations upon which the estimation of the $Q_{W}$ and $\bar{Q}$ features of $P_{0}$.

### 8.4.2. Elementary statistical properties

This subsection is very similar to its counterpart for the IPTW estimator, see Section 8.2.2.
Let us denote by $\bar{Q}_{n, 1}$ the output of algorithm $\widehat{\mathscr{A}}_{\bar{Q}, 1}$ trained on $P_{n}$, and recall that $\bar{Q}_{n, \mathrm{kNN}}$ is the output of algorithm $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$ trained on $P_{n}$. Let $\psi_{n}^{d}$ and $\psi_{n}^{e}$ be the G-computation estimators obtained by substituting $\bar{Q}_{n, 1}$ and $\bar{Q}_{n, \mathrm{kNN}}$ for $\bar{Q}_{n}$ in (31), respectively.

If $\bar{Q}_{n, \bullet}$ minimized the empirical risk over a finite-dimensional, identifiable, and well-specified working model, then $\sqrt{n}\left(\psi_{n}^{\bullet}-\psi_{0}\right)$ would converge in law to a centered Gaussian law (here $\psi_{n}^{\bullet}$ represents the G-computation estimator obtained by substituting $\bar{Q}_{n, \bullet}$ for $\bar{Q}_{n}$ in (31)). Moreover, the asymptotic variance of $\sqrt{n}\left(\psi_{n}^{\bullet}-\psi_{0}\right)$ would be estimated anti-conservatively ${ }^{19}$ with

$$
\begin{align*}
v_{n}^{d} & \doteq \operatorname{Var}_{P_{n}}\left(\bar{Q}_{n, 1}(1, \cdot)-\bar{Q}_{n, 1}(0, \cdot)\right)  \tag{32}\\
& =\frac{1}{n} \sum_{i=1}^{n}\left(\bar{Q}_{n, 1}\left(1, W_{i}\right)-\bar{Q}_{n, 1}\left(0, W_{i}\right)-\psi_{n}^{d}\right)^{2} \tag{33}
\end{align*}
$$

Unfortunately, algorithm $\widehat{\mathscr{A}}_{\bar{Q}, 1}$ is mis-specified and $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$ is not based on a finitedimensional working model. Nevertheless, function compute_gcomp estimates (in general, very poorly) the asymptotic variance with (33).

We investigate empirically the statistical behavior of $\psi_{n}^{d}$ in Section 8.4.3. For an analysis of the reason why $v_{n}^{d}$ is an anti-conservative estimator of the asymptotic variance of $\sqrt{n}\left(\psi_{n}^{d}-\psi_{0}\right)$, see there in Appendix C.1.2. We wish to emphasize that anti-conservativeness is even more embarrassing that conservativeness (both being contingent on the fact that the algorithms are well-specified, fact that cannot be true in the present case in real world situations).

What would happen if we used a less amenable algorithm $\widehat{\mathscr{A}_{\bar{Q}}}$. For instance, $\widehat{\mathscr{A}} \overline{\bar{Q}}$ could still be well-specified but so versatile/complex (as opposed to being based on well-behaved, finitedimensional parametric model) that the estimator $\bar{Q}_{n}$, though still consistent, would converge

[^13]slowly to its target. Then, root- $n$ consistency would fail to hold. We can explore empirically this situation with estimator $\psi_{n}^{e}$ that hinges on algorithm $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$. Or $\widehat{\mathscr{A}}_{\bar{Q}}$ could be mis-specified and there would be no guarantee at all that the resulting estimator $\psi_{n}$ be even consistent.

### 8.4.3. Empirical investigation, fixed sample size

Let us compute $\psi_{n}^{d}$ and $\psi_{n}^{e}$ on the same iter $=1000$ independent samples of independent observations drawn from $P_{0}$ as in Sections 6.3 and 8.2.3. We first enrich object learned_features_fixed_sample_size that was created in Section 8.2.3, adding to it estimators of $\bar{Q}_{0}$ obtained by training algorithms $\widehat{\mathscr{A}}_{\bar{Q}, 1}$ and $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$ on each smaller data set.

The second series of commands creates object psi_hat_de, an 1000 by six tibble containing notably the values of $\psi_{n}^{d}$ and $\sqrt{v_{n}^{d}} / \sqrt{n}$ computed by calling compute_gcomp, and those of the recentered (with respect to $\psi_{0}$ ) and renormalized $\sqrt{n} / \sqrt{v_{n}^{d}}\left(\psi_{n}^{d}-\psi_{0}\right)$. Because we know beforehand that $v_{n}^{d}$ under-estimates the actual asymptotic variance of $\sqrt{n}\left(\psi_{n}^{d}-\psi_{0}\right)$, the tibble also includes the values of $\sqrt{n} / \sqrt{v^{d *}}\left(\psi_{n}^{d}-\psi_{0}\right)$ where the estimator $v^{d *}$ of the asymptotic variance is computed across the replications of $\psi_{n}^{d}$. The tibble includes the same quantities pertaining to $\psi_{n}^{e}$, although there is no theoretical guarantee that the central limit theorem does hold and, even if it did, that the counterpart $v_{n}^{e}$ to $v_{n}^{d}$ estimates in any way the asymptotic variance of $\sqrt{n}\left(\psi_{n}^{e}-\psi_{0}\right)$.

Finally, bias_de reports amounts of bias (at the renormalized scales - plural). There is one value of bias for each combination of (i) type of the estimator (d or e) and (ii) how the renormalization is carried out, either based on $v_{n}^{d}$ and $v_{n}^{e}$ (auto_renormalization is TRUE) or on the estimator of the asymptotic variance computed across the replications of $\psi_{n}^{d}$ and $\psi_{n}^{e}$ (auto_renormalization is FALSE).

```
learned_features_fixed_sample_size <-
    learned_features_fixed_sample_size %>%
    mutate(Qbar_hat_d =
                map(obs,
                            ~ estimate_Qbar(., algorithm = working_model_Q_one)),
            Qbar_hat_e =
                map(obs,
                ~ estimate_Qbar(., algorithm = kknn_algo,
                    trControl = kknn_control,
                            tuneGrid = kknn_grid))) %>%
    mutate(QW = map(obs, estimate_QW),
            est_d =
                pmap(list(QW, Qbar_hat_d, n()),
                            ~ compute_gcomp(..1, wrapper(..2, FALSE), ..3)),
            est_e =
                pmap(list(QW, Qbar_hat_e, n()),
                        ~ compute_gcomp(..1, wrapper(..2, FALSE), ..3)))
psi_hat_de <- learned_features_fixed_sample_size %>%
    select(est_d, est_e) %>%
    pivot_longer(c('est_d', 'est_e'),
                            names_to = "type", values_to = "estimates") %>%
    extract(type, "type", "_([de])$") %>%
```

```
    unnest(estimates) %>%
    group_by(type) %>%
    mutate(sig_alt = sd(psi_n)) %>%
    mutate(clt_ = (psi_n - psi_zero) / sig_n,
        clt_alt = (psi_n - psi_zero) / sig_alt) %>%
    pivot_longer(c('clt_', 'clt_alt'),
        names_to = "key", values_to = "clt") %>%
    extract(key, "key", "_(.*)$") %>%
    mutate(key = ifelse(key == "", TRUE, FALSE)) %>%
    rename("auto_renormalization" = key)
(bias_de <- psi_hat_de %>%
    group_by(type, auto_renormalization) %>%
    summarize(bias = mean(clt)) %>% ungroup)
# A tibble: 4 x 3
#> type auto_renormalization bias
#> <chr> <lgl> <dbl>
#> 1 d FALSE 0.240
#> 2 dRUE 2.25
#> 3 e FALSE 0.122
#> 4 e TRUE 0.679
fig <- ggplot() +
    geom_line(aes(x = x, y = y),
        data = tibble(x = seq(-4, 4, length.out = 1e3),
                                y = dnorm(x)),
            linetype = 1, alpha = 0.5) +
    geom_density(aes(clt, fill = type, colour = type),
                psi_hat_de, alpha = 0.1) +
    geom_vline(aes(xintercept = bias, colour = type),
            bias_de, size = 1.5, alpha = 0.5) +
    facet_wrap(~ auto_renormalization,
                    labeller =
                        as_labeller(c('TRUE' = "auto-renormalization: TRUE",
                        'FALSE' = "auto-renormalization: FALSE")),
            scales = "free")
fig +
    labs(y = "",
        x = bquote(paste(sqrt(n/v[n]^{list(d, e)})*
                        (psi[n]^{list(d, e)} - psi[0]))))
```

We represent the empirical laws of the recentered (with respect to $\psi_{0}$ ) and renormalized $\psi_{n}^{d}$ and $\psi_{n}^{e}$ in Figure 11 (kernel density estimators). Two renormalization schemes are considered, either based on an estimator of the asymptotic variance (left) or on the empirical variance computed across the iter independent replications of the estimators (right). We emphasize that the $x$-axis ranges differ starkly between the left and right plots.

Two important comments are in order. First, on the one hand, the G-computation estimator


Figure 11. Kernel density estimators of the law of two G-computation estimators of $\psi_{0}$ (recentered with respect to $\psi_{0}$, and renormalized). The estimators respectively hinge on algorithms $\widehat{\mathscr{A}}_{\bar{Q}, 1}(d)$ and $\widehat{\mathscr{A}_{\bar{Q}}^{2}}, k N N(e)$ to estimate $\bar{Q}_{0}$. Two renormalization schemes are considered, either based on an estimator of the asymptotic variance (right) or on the empirical variance computed across the iter independent replications of the estimators (left). We emphasize that the $x$-axis ranges differ starkly between the left and right plots.
$\psi_{n}^{d}$ is biased. Specifically, by the above chunk of code, the averages of $\sqrt{n / v_{n}^{d}}\left(\psi_{n}^{d}-\psi_{0}\right)$ and $\sqrt{n / v_{n}^{d *}}\left(\psi_{n}^{d}-\psi_{0}\right)$ computed across the realizations are equal to 2.247 and 0.24 (see bias_de). On the other hand, the G-computation estimator $\psi_{n}^{e}$ is biased too, though slightly less than $\psi_{n}^{d}$. Specifically, by the above chunk of code, the averages of $\sqrt{n / v_{n}^{e}}\left(\psi_{n}^{e}-\psi_{0}\right)$ and $\sqrt{n / v^{e *}}\left(\psi_{n}^{e}-\psi_{0}\right)$ computed across the realizations are equal to 0.679 and 0.122 (see bias_de). We can provide an oracular explanation. Estimator $\psi_{n}^{d}$ suffers from the poor approximation of $\bar{Q}_{0}$ by $\widehat{\mathscr{A}}_{\bar{Q}, 1}\left(P_{n}\right)$, a result of the algorithm's mis-specification. As for $\psi_{n}^{e}$, it behaves better because $\widehat{\mathscr{A}} \overline{\bar{Q}}, \mathrm{kNN}\left(P_{n}\right)$ approximates $\bar{Q}_{0}$ better than $\widehat{\mathscr{A}_{\bar{Q}}^{Q}, 1}\left(P_{n}\right)$, an apparent consequence of the greater versatility of the algorithm.

Second, we get a visual confirmation that $v_{n}^{d}$ under-estimates the actual asymptotic variance of $\sqrt{n}\left(\psi_{n}^{d}-\psi_{0}\right)$ : the right-hand side red bell-shaped curve is too dispersed. In contrast, the righthand side blue bell-shaped curve is much closer to the black curve that represents the density of the standard normal law. Looking at the left-hand side plot reveals that the empirical law of $\sqrt{n / v^{d *}}\left(\psi_{n}^{d}-\psi_{0}\right)$, once translated to compensate for the bias, is rather close to the black curve. This means that the random variable is approximately distributed like a Gaussian random variable. On the contrary, the empirical law of $\sqrt{n / v^{e *}}\left(\psi_{n}^{e}-\psi_{0}\right)$ does not strike us as being as closely Gaussian-like as that of $\sqrt{n / v^{d *}}\left(\psi_{n}^{d}-\psi_{0}\right)$. By being more data-adaptive than $\widehat{\mathscr{A}}_{\bar{Q}, 1}$, algorithm $\widehat{\mathscr{A}} \overline{\bar{Q}}$,kNN yields a better estimator of $\bar{Q}_{0}$. However, the rate of convergence of $\widehat{\mathscr{A}}, \underline{\bar{Q}}, \mathrm{kNN}, ~\left(P_{n}\right)$ to its limit may be slower than root- $n$, invalidating a central limit theorem.

How do the estimated variances of $\psi_{n}^{d}$ and $\psi_{n}^{e}$ compare with their empirical counterparts (computed across the iter replications of the construction of the two estimators)?

```
## psi_n^d
(psi_hat_de %>% ungroup %>%
    filter(type == "d" & auto_renormalization) %>% pull(sig_n) %>% summary)
#> Min. 1st Qu. Median Mean 3rd Qu. Max.
#> 0.00061 0.00171 0.00208 0.00208 0.00244 0.00373
## psi_n`e
(psi_hat_de %>% ungroup %>%
    filter(type == "e" & auto_renormalization) %>% pull(sig_n) %>% summary)
#> Min. 1st Qu. Median Mean 3rd Qu. Max.
#> 0.00151 0.00261 0.00288 0.00289 0.00316 0.00430
```

The empirical standard deviation of $\psi_{n}^{d}$ is approximately 8.56 times larger than the average estimated standard deviation. The estimator is anti-conservative indeed!

As for the empirical standard deviation of $\psi_{n}^{e}$, it is approximately 6.093 times larger than the average estimated standard deviation.

### 8.4.4. 2 Empirical investigation, varying sample size

```
sample_size <- c(5e3, 15e3)
block_size <- sum(sample_size)
learned_features_varying_sample_size <- obs %>% as_tibble %>%
    head(n = (nrow(.) %/% block_size) * block_size) %>%
    mutate(block = label(1:nrow(.), sample_size)) %>%
    nest(obs = c(W, A, Y))
```

First, we cut the data set into independent sub-data sets of sample size $n$ in $\left\{5000,1.5 \times 10^{4}\right\}$. Second, we infer $\psi_{0}$ as shown two chunks earlier. We thus obtain 50 independent realizations of each estimator derived on data sets of 2 , increasing sample sizes.

```
learned_features_varying_sample_size <-
    learned_features_varying_sample_size %>%
    mutate(Qbar_hat_d =
                map(obs,
                ~ estimate_Qbar(., algorithm = working_model_Q_one)),
            Qbar_hat_e =
                map(obs,
                ~ estimate_Qbar(., algorithm = kknn_algo,
                    trControl = kknn_control,
                        tuneGrid = kknn_grid))) %>%
    mutate(QW = map(obs, estimate_QW),
                est_d =
                pmap(list(QW, Qbar_hat_d, n()),
                            ~ compute_gcomp(..1, wrapper(..2, FALSE), ..3)),
            est_e =
                pmap(list(QW, Qbar_hat_e, n()),
                        ~ compute_gcomp(..1, wrapper(..2, FALSE), ..3)))
```

```
root_n_bias <- learned_features_varying_sample_size %>%
    mutate(block = unlist(map(strsplit(block, "_"), ~.x[2])),
            sample_size = sample_size[as.integer(block)]) %>%
    select(block, sample_size, est_d, est_e) %>%
    pivot_longer(c('est_d', 'est_e'),
                            names_to = "type", values_to = "estimates") %>%
    extract(type, "type", "_([de])$") %>%
    unnest(estimates) %>%
    group_by(block, type) %>%
    mutate(sig_alt = sd(psi_n)) %>%
    mutate(clt_ = (psi_n - psi_zero) / sig_n,
            clt_alt = (psi_n - psi_zero) / sig_alt) %>%
    pivot_longer(c('clt_', 'clt_alt'),
                            names_to = "key", values_to = "clt") %>%
    extract(key, "key", "_(.*)$") %>%
    mutate(key = ifelse(key == "", TRUE, FALSE)) %>%
    rename("auto_renormalization" = key)
```

The tibble called root_n_bias reports root- $n$ times bias for all combinations of estimator and sample size. The next chunk of code presents visually our findings, see Figure 12. Note how we include the realizations of the estimators derived earlier and contained in psi_hat_de (thus breaking the independence between components of root_n_bias, a small price to pay in this context).

```
root_n_bias <- learned_features_fixed_sample_size %>%
    mutate(block = "0",
                sample_size = B/iter) %>% # because *fixed* sample size
    select(block, sample_size, est_d, est_e) %>%
    pivot_longer(c('est_d', 'est_e'),
                            names_to = "type", values_to = "estimates") %>%
    extract(type, "type", "_([de])$") %>%
    unnest(estimates) %>%
    group_by(block, type) %>%
    mutate(sig_alt = sd(psi_n)) %>%
    mutate(clt_ = (psi_n - psi_zero) / sig_n,
                clt_alt = (psi_n - psi_zero) / sig_alt) %>%
    pivot_longer(c('clt_', 'clt_alt'),
                names_to = "key", values_to = "clt") %>%
    extract(key, "key", "_(.*)$") %>%
    mutate(key = ifelse(key == "", TRUE, FALSE)) %>%
    rename("auto_renormalization" = key) %>%
    full_join(root_n_bias)
root_n_bias %>% filter(auto_renormalization) %>%
    mutate(rnb = sqrt(sample_size) * (psi_n - psi_zero)) %>%
    group_by(sample_size, type) %>%
    ggplot() +
    stat_summary(aes(x = sample_size, y = rnb,
```

```
    group = interaction(sample_size, type),
        color = type),
    fun.data = mean_se, fun.args = list(mult = 1),
    position = position_dodge(width = 250), cex = 1) +
stat_summary(aes(x = sample_size, y = rnb,
            group = interaction(sample_size, type),
            color = type),
    fun.data = mean_se, fun.args = list(mult = 1),
    position = position_dodge(width = 250), cex = 1,
    geom = "errorbar", width = 750) +
stat_summary(aes(x = sample_size, y = rnb,
            color = type),
        fun = mean,
        position = position_dodge(width = 250),
        geom = "polygon", fill = NA) +
geom_point(aes(x = sample_size, y = rnb,
            group = interaction(sample_size, type),
            color = type),
        position = position_dodge(width = 250),
        alpha = 0.1) +
scale_x_continuous(breaks = unique(c(B / iter, sample_size))) +
labs(x = "sample size n",
    y = bquote(paste(sqrt(n) * (psi[n]^{list(d, e)} - psi[0]))))
```



FIGURE 12. Evolution of root-n times bias versus sample size for two $G$-computation estimators of $\psi_{0}$. The estimators respectively hinge on algorithms $\widehat{\mathscr{A}}_{\bar{Q}}^{1}, 1(d)$ and $\widehat{\mathscr{A}}_{\bar{Q}}, k N N$ (e) to estimate $\bar{Q}_{0}$. Big dots represent the average biases and vertical lines represent twice the standard error.

Root- $n$ bias for $\psi_{n}^{d}$ (red lines and points) is positive and tends to increase from sample size 1000 to 5000 and from 5000 to $1.5 \times 10^{4}$. Moreover, root- $n$ bias tends to be larger for $\psi_{n}^{d}$ than for $\psi_{n}^{e}$. In addition, root- $n$ bias for $\psi_{n}^{e}$ tends to increase from sample size 1000 to $1.5 \times 10^{4}$, where it
tends to be positive too. ${ }^{20}$
In essence, we observe that the bias does not vanish faster than root- $n$. For $\psi_{n}^{d}$, this is because $\widehat{\mathscr{A}}_{\bar{Q}, 1}$ is mis-specified and we expect that the bias increases at rate root- $n$. For $\psi_{n}^{e}$, this is because the estimator produced by the versatile $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$ converges slowly to its limit. Can anything be done to amend $\psi_{n}^{d}$ and $\psi_{n}^{e}$ ?

### 8.5. Investigating further the G-computation estimation strategy

1. Implement the G-computation estimator based on well-specified working model. To do so, (i) create a copy working_model_Q_two of working_model_Q_one, (ii) replace its family entry in such a way that $\bar{Q}_{0}$ falls in the corresponding working model, and (iii) adapt the chunk of code where we computed Qbar_hat_d in Section 8.4.1.
2. Evaluate the empirical properties of the estimator of problem 1.
3. Create a function to estimate the marginal law $Q_{0, W}$ of $W$ by maximum likelihood estimation based on a mis-specified model that assumes (wrongly) that $W$ is drawn from a Gaussian law with unknown mean and variance.
4. 2 Implement the G-computation estimator using either working_model_Q_one or working_model_Q_two and the estimator of $Q_{0, W}$ of problem 3 .

## 9. One-step correction

### 9.1. Z General analysis of plug-in estimators

Recall that $\widehat{\mathscr{A}}_{Q_{W}}$ is an algorithm designed for the estimation of $Q_{0, W}$ (see Section 7.3) and that we denote by $Q_{n, W} \doteq \widehat{\mathscr{A}}_{Q_{W}}\left(P_{n}\right)$ the output of the algorithm trained on $P_{n}$. Likewise, $\widehat{\mathscr{A}}_{\bar{G}}$ and $\widehat{\mathscr{A}}_{\bar{Q}}$ are two generic algorithms designed for the estimation of $\bar{G}_{0}$ and of $\bar{Q}_{0}$ (see Sections 7.4 and 7.6), $\bar{G}_{n} \doteq \widehat{\mathscr{A}} \overline{\bar{G}}\left(P_{n}\right)$ and $\bar{Q}_{n} \doteq \widehat{\mathscr{A}}_{\bar{Q}}\left(P_{n}\right)$ are their outputs once trained on $P_{n}$.

Let us now introduce $P_{n}^{\circ}$ a law in $\mathscr{M}$ such that the $Q_{W}, \bar{G}$ and $\bar{Q}$ features of $P_{n}^{\circ}$ equal $Q_{n, W}, \bar{G}_{n}$ and $\bar{Q}_{n}$, respectively. We say that any such law is compatible with $Q_{n, W}, \bar{G}_{n}$ and $\bar{Q}_{n}$.

Now, let us substitute $P_{n}^{\circ}$ for $P$ in (19):

$$
\begin{equation*}
\Psi\left(P_{n}^{\circ}\right)-\Psi\left(P_{0}\right)=-P_{0} D^{*}\left(P_{n}^{\circ}\right)+\operatorname{Rem}_{P_{0}}\left(P_{n}^{\circ}\right) \tag{34}
\end{equation*}
$$

We show there in Appendix C. 2 that, under conditions on the complexity/versatility of algorithms $\widehat{\mathscr{A}}_{\bar{G}}$ and $\widehat{\mathscr{A}}_{\bar{Q}}$ (often referred to as regularity conditions) and assuming that their outputs $\bar{G}_{n}$ and $\bar{Q}_{n}$ both consistently estimate their targets $\bar{G}_{0}$ and $\bar{Q}_{0}$, it holds that

$$
\begin{align*}
\Psi\left(P_{n}^{\circ}\right)-\Psi\left(P_{0}\right) & =-P_{n} D^{*}\left(P_{n}^{\circ}\right)+P_{n} D^{*}\left(P_{0}\right)+o_{P_{0}}(1 / \sqrt{n})  \tag{35}\\
& =-P_{n} D^{*}\left(P_{n}^{\circ}\right)+\frac{1}{n} \sum_{i=1}^{n} D^{*}\left(P_{0}\right)\left(O_{i}\right)+o_{P_{0}}(1 / \sqrt{n}) .
\end{align*}
$$

[^14]In light of (16), $\Psi\left(P_{n}^{\circ}\right)$ would be asymptotically linear with influence curve $\mathrm{IC}=D^{*}\left(P_{0}\right)$ in the absence of the random term $-P_{n} D^{*}\left(P_{n}^{\circ}\right)$. Unfortunately, it turns out that this term can degrade dramatically the behavior of the plug-in estimator $\Psi\left(P_{n}^{\circ}\right)$.

### 9.2. One-step correction

Luckily, a very simple workaround allows to circumvent the problem. Proposed by Le Cam (1969) (see also (Pfanzagl, 1982) and (van der Vaart, 1998)), the workaround merely consists in adding the random term to the initial estimator, that is, in estimating $\Psi\left(P_{0}\right)$ not with $\Psi\left(P_{n}^{\circ}\right)$ but instead with

$$
\begin{equation*}
\psi_{n}^{\mathrm{os}} \doteq \Psi\left(P_{n}^{\circ}\right)+P_{n} D^{*}\left(P_{n}^{\circ}\right)=\Psi\left(P_{n}^{\circ}\right)+\frac{1}{n} \sum_{i=1}^{n} D^{*}\left(P_{n}^{\circ}\right)\left(O_{i}\right) \tag{36}
\end{equation*}
$$

Obviously, in light of (35), $\psi_{n}^{\text {os }}$ is asymptotically linear with influence curve IC $=D^{*}\left(P_{0}\right)$. Thus, by the central limit theorem, $\sqrt{n}\left(\psi_{n}^{\text {os }}-\Psi\left(P_{0}\right)\right)$ converges in law to a centered Gaussian distribution with variance

$$
\begin{equation*}
\operatorname{Var}_{P_{0}}\left(D^{*}\left(P_{0}\right)(O)\right)=\mathrm{E}_{P_{0}}\left(D^{*}\left(P_{0}\right)(O)\right) \tag{37}
\end{equation*}
$$

The detailed general analysis of plug-in estimators developed there in Appendix C. 2 also revealed that the above asymptotic variance is consistently estimated with

$$
\begin{equation*}
P_{n} D^{*}\left(P_{n}^{\circ}\right)^{2}=\frac{1}{n} \sum_{i=1}^{n} D^{*}\left(P_{n}^{\circ}\right)^{2}\left(O_{i}\right) \tag{38}
\end{equation*}
$$

Therefore, by the central limit theorem and Slutsky's lemma (see the argument there in Appendix B.3.1),

$$
\left[\psi_{n}^{\mathrm{os}} \pm \Phi^{-1}(1-\alpha) \frac{P_{n} D^{*}\left(P_{n}^{\circ}\right)^{2}}{\sqrt{n}}\right]
$$

is a confidence interval for $\Psi\left(P_{0}\right)$ with asymptotic level $(1-2 \alpha)$.

### 9.3. Empirical investigation

In light of (36) if the estimator equals $\Psi\left(P_{n}^{\circ}\right)$, then performing a one-step correction essentially boils down to computing two quantities, $-P_{n} D^{*}\left(P_{n}^{\circ}\right)$ (the bias term) and $P_{n} D^{*}\left(P_{n}^{\circ}\right)^{2}$ (an estimator of the asymptotic variance of $\left.\psi_{n}^{\mathrm{os}}\right)$. The tlrider package makes the operation very easy thanks to the function apply_one_step_correction.

Let us illustrate its use by updating the G-computation estimator built on the $n=1000$ first observations in obs by relying on $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$, that is, on the algorithm for the estimation of $\bar{Q}_{0}$ as it is implemented in estimate_Qbar with its argument algorithm set to the built-in kknn_algo (see Section 7.6.2). The algorithm has been trained earlier on this data set and produced the object Qbar_hat_kknn. The following chunk of code re-computes the corresponding G-computation estimator, using again the estimator QW_hat of the marginal law of $W$ under $P_{0}$ (see Section 7.3), then applied the one-step correction:

```
(psin_kknn <- compute_gcomp(QW_hat, wrapper(Qbar_hat_kknn, FALSE), 1e3))
#> # A tibble: 1 x 2
#> psi_n sig_n
#> <dbl> <dbl>
#> 1 0.0887 0.00245
(psin_kknn_os <- apply_one_step_correction(head(obs, 1e3),
                                    wrapper(Gbar_hat, FALSE),
                                    wrapper(Qbar_hat_kknn, FALSE),
                                    psin_kknn$psi_n))
#> # A tibble: 1 x 3
#> psi_n sig_n crit_n
#> <dbl> <dbl> <db\imath`>
#> 1 0.0888 0.0162 0.000124
```

In the call to apply_one_step_correction we provide (i) the data set at hand (first line), (ii) the estimator Gbar_hat of $\bar{G}_{0}$ that we built earlier by using algorithm $\widehat{\mathscr{A}}_{\bar{G}, 1}$ (second line; see Section 7.4.2), (iii) the estimator Qbar_hat_kknn of $\bar{Q}_{0}$ and the G-computation estimator psin_kknn that resulted from it (third and fourth lines).

To assess what is the impact of the one-step correction, let us apply the one-step correction to the estimators that we built in Section 8.4.3. The object learned_features_fixed_sample_size already contains the estimated features of $P_{0}$ that are needed to perform the one-step correction to the estimators $\psi_{n}^{d}$ and $\psi_{n}^{e}$, namely, thus we merely have to call the function apply_one_step_correction.

```
psi_hat_de_one_step <- learned_features_fixed_sample_size %>%
    mutate(os_est_d =
            pmap(list(obs, Gbar_hat, Qbar_hat_d, est_d),
                            ~ apply_one_step_correction(as.matrix(..1),
                                    wrapper(..2, FALSE),
                                    wrapper(..3, FALSE),
                                    ..4$psi_n)),
            os_est_e =
                        pmap(list(obs, Gbar_hat, Qbar_hat_e, est_e),
                            ~ apply_one_step_correction(as.matrix(..1),
                                    wrapper(..2, FALSE),
                                    wrapper(..3, FALSE),
                                    ..4$psi_n))) %>%
    select(os_est_d, os_est_e) %>%
    pivot_longer(c('os_est_d', 'os_est_e'),
                            names_to = "type", values_to = "estimates") %>%
    extract(type, "type", "_([de])$") %>%
    mutate(type = pasteO(type, "_one_step")) %>%
    unnest(estimates) %>%
    group_by(type) %>%
    mutate(sig_alt = sd(psi_n)) %>%
    mutate(clt_ = (psi_n - psi_zero) / sig_n,
            clt_alt = (psi_n - psi_zero) / sig_alt) %>%
    pivot_longer(c('clt_', 'clt_alt'),
```

```
        names_to = "key", values_to = "clt") %>%
    extract(key, "key", "_(.*)$") %>%
    mutate(key = ifelse(key == "", TRUE, FALSE)) %>%
    rename("auto_renormalization" = key)
(bias_de_one_step <- psi_hat_de_one_step %>%
    group_by(type, auto_renormalization) %>%
    summarize(bias = mean(clt)) %>% ungroup)
#> # A tibble: 4 x 3
#> type auto_renormalization bias
#> <chr> <lgl> <dbl>
#> 1 d_one_step FALSE -0.00668
#> 2 d_one_step TRUE -0.0240
#> 3 e_one_step FALSE 0.0333
#> 4 e_one_step TRUE 0.0171
fig <- ggplot() +
    geom_line(aes(x = x, y = y),
                data = tibble(x = seq(-4, 4, length.out = 1e3),
                    y = dnorm(x)),
                linetype = 1, alpha = 0.5) +
    geom_density(aes(clt, fill = type, colour = type),
                        psi_hat_de_one_step, alpha = 0.1) +
    geom_vline(aes(xintercept = bias, colour = type),
                            bias_de_one_step, size = 1.5, alpha = 0.5) +
    facet_wrap(~ auto_renormalization,
            labeller =
                        as_labeller(c('TRUE' = "auto-renormalization: TRUE",
                            'FALSE' = "auto-renormalization: FALSE")),
            scales = "free")
fig +
    labs(y = "",
        x = bquote(paste(sqrt(n/v[n]^{list(d, e, os)})*
                            (psi[n]~{list(d, e, os)} - psi[0]))))
```

It seems that the one-step correction performs qui well (in particular, compare bias_de with bias_de_one_step):

```
bind_rows(bias_de, bias_de_one_step) %>%
    filter(!auto_renormalization) %>%
    arrange(type)
#> # A tibble: 4 x 3
#> type auto_renormalization bias
#> <chr> <lgl> <dbl>
#> 1 d FALSE 0.240
#> 2 d_one_step FALSE -0.00668
#> 3 e FALSE 0.122
#> 4 e_one_step FALSE 0.0333
```



Figure 13. Kernel density estimators of the law of two one-step-G-computation estimators of $\psi_{0}$ (recentered with respect to $\psi_{0}$, and renormalized). The estimators respectively hinge on algorithms $\widehat{\mathscr{A}}_{\bar{Q}, 1}(d)$ and $\widehat{\mathscr{A}}_{\overline{\mathrm{Q}}}^{, k N N}, ~(e)$ to estimate $\bar{Q}_{0}$, and on one-step correction. Two renormalization schemes are considered, either based on an estimator of the asymptotic variance (left) or on the empirical variance computed across the iter independent replications of the estimators (right). We emphasize that the $x$-axis ranges differ between the left and right plots.

What about the estimation of the asymptotic variance, and of the mean squared-error of the estimators?

```
psi_hat_de %>%
    full_join(psi_hat_de_one_step) %>%
    filter(auto_renormalization) %>%
    group_by(type) %>%
    summarize(sd = mean(sig_n),
            se = sd(psi_n),
            mse = mean((psi_n - psi_zero)^2) * n()) %>%
    arrange(type)
#> # A tibble: 4 x 4
#> type sd se mse
#> <chr> <dbl> <dbl> <dbl>
#> 1 d 0.00208 0.0178 0.336
#> 2 d_one_step 0.0174 0.0174 0.302
#> 3 e 0.00289 0.0176 0.314
#> 4 e_one_step 0.0167 0.0176 0.309
```

The sd (estimator of the asymptotic standard deviation) and se (empirical standard deviation) entries are similar. This indicates that the inference of the asymptotic variance of the one-step estimators based on the influence curve is rather accurate for both the d-and e-variants that we implemented. As for the mean square error, it is diminished by the one-step update for both types $d$ and e, the e_one_step estimator exhibiting the smallest mean square error.

### 9.4. Investigating further the one-step correction methodology

1. Use estimate_Gbar to create an oracle algorithm $\widetilde{\mathscr{A}}_{\bar{G}, S}$ for the estimation of $\bar{G}_{0}$ that, for any $s>0$, estimates $\bar{G}_{0}(w)$ with

$$
\bar{G}_{n}(w) \doteq \widetilde{\mathscr{A}}_{\bar{G}, s}\left(P_{n}\right)(w) \doteq \operatorname{expit}\left(\operatorname{logit}\left(\bar{G}_{0}(w)\right)+s Z\right)
$$

where $Z$ is a standard normal random variable. ${ }^{21}$ What would happen if one chose $s=0$ in the above definition? What happens when $s$ converges to 0 ? Explain why the algorithm is said to be an oracle algorithm.
2. Use estimate_Qbar to create an oracle algorithm $\widetilde{\mathscr{A}_{\bar{Q}, s}}$ for the estimation of $\bar{Q}_{0}$ that, for any $s>0$, estimates $\bar{Q}_{0}(a, w)$ with

$$
\bar{Q}_{n}(a, w) \doteq \widetilde{\mathscr{A}}_{\bar{G}, s}\left(P_{n}\right)(a, w) \doteq \operatorname{expit}\left(\operatorname{logit}\left(\bar{Q}_{0}(a, w)\right)+s Z\right)
$$

where $Z$ is a standard normal random variable. The comments made about $\widetilde{\mathscr{A}}_{\bar{G}, s}$ in the above problem also apply to $\widetilde{\mathscr{A}_{\bar{Q}}, s}$.
3. Reproduce the simulation study developed in Sections 8.4 .3 and 9.3 with the oracle algorithms $\widetilde{\mathscr{A}}_{\bar{G}, s}$ and $\widetilde{\mathscr{A}}_{\bar{Q}, s^{\prime}}$ substituted for used in these sections. Change the values of $s, s^{\prime}>0$ and compare how well the estimating procedure performs depending on the product $s s^{\prime}$. What do you observe? We invite you to refer to Section C.2.1.

## 10. Targeted minimum loss-based estimation

### 10.1. Motivations

### 10.1.1. Falling outside the parameter space

Section 9 introduced the one-step corrected estimator $\psi_{n}^{\text {os }}$ of $\psi_{0}$. It is obtained by adding a correction term to an initial plug-in estimator $\Psi\left(P_{n}^{\circ}\right)$, resulting in an estimator that is asymptotically linear with influence curve IC $=D^{*}\left(P_{0}\right)$ under mild conditions (see Section 9.2 and the detailed proof there in Appendix C.2).

Unfortunately, the one-step estimator lacks a desirable feature of a plug-in estimator: plug-in estimators always lie in the parameter space whereas a one-step estimator does not necessarily do so. For example, it must also be true that $\psi_{0}=\Psi\left(P_{0}\right) \in[-1,1]$ yet it may be the case that $\psi_{n}^{\text {os }} \notin[-1,1]$.

It is typically easy to shape an algorithm $\widehat{\mathscr{A}}_{\bar{Q}}$ for the estimation of $\bar{Q}_{0}$ to output estimators $\bar{Q}_{n}$ that, like $\bar{Q}_{0}$, take their values in $[0,1]$. The plug-in estimator $\psi_{n}(31)$ based on such a $\bar{Q}_{n}$ necessarily satisfies $\psi_{n} \in[-1,1]$. However, the one-step estimator derived from $\psi_{n}$ may fall outside of the interval if the correction term $P_{n} D^{*}\left(P_{n}^{\circ}\right)$ is large.

[^15]
### 10.1.2. The influence curve equation

Upon closer examination of the influence curve $D^{*}\left(P_{n}^{\circ}\right)$, we see that this may occur more frequently when $\ell \bar{G}_{n}\left(A_{i}, W_{i}\right)$ is close to zero for at least some $1 \leq i \leq n$. In words, this may happen if there are observations in our data set that we observed under actions $A_{i}$ that were estimated to be unlikely given their respective contexts $W_{i}$. In such cases, $D^{*}\left(P_{n}^{\circ}\right)\left(O_{i}\right)$, and consequently the one-step correction term, may be large and cause the one-step estimate to fall outside $[-1,1]$.

Another way to understand this behavior is to recognize the one-step estimator as an initial plug-in estimator that is corrected in the parameter space of $\psi_{0}$. One of the pillars of targeted learning is to perform, instead, a correction in the parameter space of $\bar{Q}_{0}$.

In particular, consider a law $P_{n}^{*}$ estimating $P_{0}$ that is compatible with the estimators $\bar{Q}_{n}^{*}, \bar{G}_{n}$, and $Q_{n, W}$, but moreover is such that

$$
\begin{equation*}
P_{n} D^{*}\left(P_{n}^{*}\right)=0 \tag{39}
\end{equation*}
$$

or, at the very least,

$$
\begin{equation*}
P_{n} D^{*}\left(P_{n}^{*}\right)=o_{P_{0}}(1 / \sqrt{n}) . \tag{40}
\end{equation*}
$$

Achieving (39) is called solving the efficient influence curve equation. Likewise, achieving (40) is called approximately solving the influence curve equation.

If such estimators can be generated indeed, then the plug-in estimator

$$
\psi_{n}^{*} \doteq \Psi\left(P_{n}^{*}\right)=\int\left(\bar{Q}_{n}^{*}(1, w)-\bar{Q}_{n}^{*}(0, w)\right) d Q_{n, W}(w)
$$

is asymptotically linear with influence curve IC $=D^{*}\left(P_{0}\right)$, under mild conditions. Moreover, by virtue of its plug-in construction, it has the additional property that in finite-samples $\psi_{n}^{*}$ will always obey bounds on the parameter space.

### 10.1.3. A basic fact on the influence curve equation

Our strategy for constructing such a plug-in estimate begins by generating an initial estimator $\bar{Q}_{n}$ of $\bar{Q}_{0}$ based on an algorithm $\widehat{\mathscr{A}} \overline{\bar{Q}}$ and an estimator $\bar{G}_{n}$ of $\bar{G}_{0}$ based on an algorithm $\widehat{\mathscr{A}_{\bar{G}}}$. These initial estimators should strive to be as close as possible to their respective targets. We use the empirical distribution $Q_{n, W}$ as an estimator of $Q_{0, W}$.

Now, recall the definition of $D_{1}^{*}$ (12) and note that for any estimator $\bar{Q}_{n}^{*}$ of $\bar{Q}_{0}$ and a law $P_{n}^{*}$ that is compatible with $\bar{Q}_{n}^{*}$ and $Q_{n, W}$,

$$
P_{n} D_{1}\left(P_{n}^{*}\right)=0 .
$$

The proof can be found there in Appendix C.4.1.
In words, this shows that so long as we use the empirical distribution $Q_{n, W}$ to estimate $Q_{0, W}$, by its very construction achieving (39) or (40) is equivalent to solving

$$
\begin{equation*}
P_{n} D_{2}^{*}\left(P_{n}^{*}\right)=0 \tag{41}
\end{equation*}
$$

or

$$
\begin{equation*}
P_{n} D_{2}^{*}\left(P_{n}^{*}\right)=o_{P_{0}}(1 / \sqrt{n}) . \tag{42}
\end{equation*}
$$

It thus remains to devise a strategy for ensuring that $P_{n} D_{2}^{*}\left(P_{n}^{*}\right)=0$ or $o_{P_{0}}(1 / \sqrt{n})$.

### 10.2. Targeted fluctuation

Our approach to satisfying (41) ${ }^{22}$ hinges on revising our initial estimator $\bar{Q}_{n}$ of $\bar{Q}_{0}$. We propose a method for building an estimator of $\bar{Q}_{0}$ that is "near to" $\bar{Q}_{n}$, but is such that for any law $P_{n}^{*}$ that is compatible with this revised estimator, $P_{n} D^{*}\left(P_{n}^{*}\right)=0 .{ }^{23}$ Because $\bar{Q}_{n}$ is our best (initial) estimator of $\bar{Q}_{0}$, the new estimator that we shall build should be at least as good an estimator of $\bar{Q}_{0}$ as $\bar{Q}_{n}$. So first, we must propose a way to move between regression functions, and then we must propose a way to move to a new regression function that fits the data at least as well as $\bar{Q}_{n}$.

Instead of writing "propos[ing] a way to move between regression functions" we may also have written "proposing a way to fluctuate a regression function", thus suggesting very opportunely that the notion of fluctuation as discussed in Section 3.3 may prove instrumental to achieve the former objective.

### 10.2.1. 2 Fluctuating indirectly

Let us resume the discussion where we left it at the end of Section 3.3.1. It is easy to show that the fluctuation $\left\{P_{h}: h \in H\right\}$ of $P$ in direction of $s$ in $\mathscr{M}$ induces a fluctuation $\left\{\bar{Q}_{h}: h \in\right.$ $H\}$ of $\bar{Q}=\left.Q_{h}\right|_{h=0}$ in the space $\mathscr{Q} \doteq\{\bar{Q}: P \in \mathscr{M}\}$ of regression functions induced by model $\mathscr{M}$. Specifically we show there in Appendix C.4.2 that, for every $h \in H$, the conditional mean $\bar{Q}_{h}(A, W)$ of $Y$ given $(A, W)$ under $P_{h}$ is given by

$$
\bar{Q}_{h}(A, W) \doteq \frac{\bar{Q}(A, W)+h \mathrm{E}_{P}(Y s(O) \mid A, W)}{1+h \mathrm{E}_{P}(s(O) \mid A, W)}
$$

We note that if $s(O)$ depends on $O$ only through ( $A, W$ ) then, abusing notation and writing $s(A, W)$ for $s(O)$,

$$
\begin{align*}
\bar{Q}_{h}(A, W) & =\frac{\bar{Q}(A, W)+h \mathrm{E}_{P}(Y s(A, W) \mid A, W)}{1+h \mathrm{E}_{P}(s(A, W) \mid A, W)} \\
& =\frac{\bar{Q}(A, W)+h s(A, W) \bar{Q}(A, W)}{1+h s(A, W)} \\
& =\bar{Q}(A, W) . \tag{43}
\end{align*}
$$

In words, $\bar{Q}$ is not fluctuated at all, that is, the laws $P_{h}$ that are elements of the fluctuation share the same conditional mean of $Y$ given $(A, W) .{ }^{24}$

However we find it easier in the present context, notably from a computational perspective, to fluctuate $\bar{Q}_{n}$ directly in $\mathscr{Q}$, as opposed to indirectly through a fluctuation defined in $\mathscr{M}$ of a law compatible with $\bar{Q}_{n}$. The next section introduces such a direct fluctuation.

### 10.2.2. Fluctuating directly

Set arbitrarily $\bar{Q} \in \mathscr{Q}$. The (direct) fluctuation of $\bar{Q}$ that we propose to consider depends on a usersupplied $\bar{G}$. For any $\bar{G}$ such that $0<\ell \bar{G}(A, W)<1, P_{0}$-almost surely, the $\bar{G}$-specific fluctuation

[^16]model for $\bar{Q}$ is
\[

$$
\begin{equation*}
\mathscr{Q}(\bar{Q}, \bar{G}) \doteq\left\{(w, a) \mapsto \bar{Q}_{h}(a, w) \doteq \operatorname{expit}\left(\operatorname{logit}(\bar{Q}(a, w))+h \frac{2 a-1}{\ell \bar{G}(a, w)}\right): t \in \mathbb{R}\right\} \subset \mathscr{Q} \tag{44}
\end{equation*}
$$

\]

Fluctuation $\mathscr{Q}(\bar{Q}, \bar{G})$ is a one-dimensional parametric model (indexed by the real-valued parameter $h$ ) that goes through $\bar{Q}$ (at $h=0$ ). For each $h \in \mathbb{R}, \bar{Q}_{h} \in \mathscr{Q}$ is the conditional mean of $Y$ given $(A, W)$ under infinitely many laws $P \in \mathscr{M}$.

The following chunk of code represents three elements of the fluctuations $\mathscr{Q}\left(\bar{Q}_{0}, \bar{G}_{0}\right)$ and $\mathscr{Q}\left(\bar{Q}_{n, \text { trees }}, \bar{G}_{0}\right)$, where $\bar{Q}_{n, \text { trees }}$ is an estimator of $\bar{Q}_{0}$ derived by the boosted trees algorithm (see Section 7.6.3).

```
Qbar_hminus <- fluctuate(Qbar, Gbar, h = -1)
Qbar_hplus <- fluctuate(Qbar, Gbar, h = +1)
Qbar_trees <- wrapper(Qbar_hat_trees, FALSE)
Qbar_trees_hminus <- fluctuate(Qbar_trees, Gbar, h = -1)
Qbar_trees_hplus <- fluctuate(Qbar_trees, Gbar, h = +1)
tibble(w = seq(0, 1, length.out = 1e3)) %>%
    mutate(truth_0_1 = Qbar(cbind(A = 1, W = w)),
        truth_0_0 = Qbar(cbind(A = 0, W = w)),
        trees_0_1 = Qbar_trees(cbind(A = 1, W = w)),
        trees_0_0 = Qbar_trees(cbind(A = 0, W = w)),
        truth_hminus_1 = Qbar_hminus(cbind (A = 1, W = w)),
        truth_hminus_0 = Qbar_hminus(cbind(A = 0, W = w)),
        trees_hminus_1 = Qbar_trees_hminus(cbind(A = 1, W = w)),
        trees_hminus_0 = Qbar_trees_hminus(cbind(A = 0, W = w)),
        truth_hplus_1 = Qbar_hplus(cbind(A = 1, W = w)),
        truth_hplus_0 = Qbar_hplus(cbind(A = 0, W = w)),
        trees_hplus_1 = Qbar_trees_hplus(cbind(A = 1, W = w)),
        trees_hplus_0 = Qbar_trees_hplus(cbind(A = 0, W = w))) %>%
    pivot_longer(-w, names_to = "f", values_to = "value") %>%
    extract(f, c("f", "h", "a"), "([^_]+)_([-_]+)_([01]+)") %>%
    mutate(f = ifelse(f == "truth", "Q_0", "Q_n"),
        h = factor(ifelse(h == "0", 0, ifelse(h == "hplus", 1, -1)))) %>%
    mutate(a = paste0("a=", a),
        fh = pasteO("(", f, ",", h, ")")) %>%
    ggplot +
    geom_line(aes(x = w, y = value, color = h, linetype = f, group = fh),
                size = 1) +
    labs(y = bquote(paste(f[h] (a,w))),
        title = bquote("Visualizing three elements of two fluctuations of"
                            ~ bar(Q)[0] ~ "and" ~ bar(Q)[n])) +
    ylim(NA, 1) +
    facet_wrap(~ a)
```



Figure 14. Representing three elements $\bar{Q}_{h}$ of the fluctuations $\mathscr{Q}\left(\bar{Q}_{0}, \bar{G}_{0}\right)$ and $\mathscr{Q}\left(\bar{Q}_{n \text {.trees }}, \bar{G}_{0}\right)$, respectively, where $\bar{Q}_{n, \text { trees }}$ is an estimator of $\bar{Q}_{0}$ derived by the boosted trees algorithm (see Section 7.6.3). The three elements correspond to $h=-1,0,1$. When $h=0, \bar{Q}_{h}$ equals either $\bar{Q}_{0}$ or $\bar{Q}_{n, t r e e s}$, depending on which fluctuation is roamed.

### 10.2.3. More on fluctuations

1. Justify the series of equalities in (43).
2. Justify that, in fact, if $s(O)$ depends on $O$ only through $(A, W)$ then the conditional law of $Y$ given $(A, W)$ under $P_{h}$ equals that under $P$.

### 10.2.4. Targeted roaming of a fluctuation

Recall that our goal is to build an estimator of $\bar{Q}_{0}$ that is at least as good as $\bar{Q}_{n}$. We will look for this enhanced estimator in the fluctuation $\mathscr{Q}\left(\bar{Q}_{n}, \bar{G}_{n}\right)$ where $\bar{G}_{n}$ is an estimator of $\bar{G}_{0}$ that is bounded away from 0 and 1 . Thus, the estimator writes as $\bar{Q}_{n, h_{n}}$ for some data-driven $h_{n} \in \mathbb{R}$. We will clarify why we do so at a later time

To assess the performance of all the candidates included in the fluctuation, we formally rely on the empirical risk function $h \mapsto \mathrm{E}_{P_{n}}\left(L_{y}\left(\bar{Q}_{n, h}\right)(O)\right)$ where

$$
\mathrm{E}_{P_{n}}\left(L_{y}\left(\bar{Q}_{n, h}\right)(O)\right)=\frac{1}{n} \sum_{i=1}^{n} L_{y}\left(\bar{Q}_{n, h}\right)\left(O_{i}\right)
$$

and the logistic (or negative binomial) loss function $L_{y}$ is given by

$$
-L_{y}(f)(O) \doteq Y \log f(A, W)+(1-Y) \log (1-f(A, W))
$$

for any function $f:\{0,1\} \times[0,1] \rightarrow[0,1]$. This loss function is the counterpart of the loss function $L_{a}$ defined in (30). The justification that we gave in Section 7.4.1 of the relevance of $L_{a}$ also applies here, mutatis mutandis. In summary, the oracle risk $\mathrm{E}_{P_{0}}\left(L_{y}(f)(O)\right)$ of $f$ is a realvalued measure of discrepancy between $\bar{Q}_{0}$ and $f ; \bar{Q}_{0}$ minimizes $f \mapsto \mathrm{E}_{P_{0}}\left(L_{y}(f)(O)\right)$ over the set of all (measurable) functions $f:[0,1] \times\{0,1\} \times[0,1] \rightarrow[0,1]$; and a minimizer $h_{0} \in \mathbb{R}$ of
$h \mapsto \mathrm{E}_{P_{0}}\left(L_{y}\left(\bar{Q}_{n, h}\right)(O)\right)$ identifies the element of fluctuation $\mathscr{Q}\left(\bar{Q}_{n}, \bar{G}_{n}\right)$ that is closest to $\bar{Q}_{0}$ (some details are given there in Appendix B.6).

It should not come as a surprise after this discussion that the aforeentioned data-driven $h_{n}$ is chosen to be the minimizer of the empirical risk function, i.e.,

$$
h_{n} \doteq \underset{h \in \mathbb{R}}{\arg \min } \mathrm{E}_{P_{n}}\left(L_{y}\left(\bar{Q}_{n, h}\right)(O)\right)
$$

The criterion is convex so the optimization program is well-posed.
The next chunk of code illustrates the search of an approximation of $h_{n}$ over a grid of candidate values.

```
candidates <- seq(-0.01, 0.01, length.out = 1e4)
W <- obs[1:1e3, "W"]
A <- obs[1:1e3, "A"]
Y <- obs[1:1e3, "Y"]
lGAW <- compute_lGbar_hatAW(A, W, Gbar_hat)
QAW <- compute_Qbar_hatAW(A, W, Qbar_hat_trees)
risk <- sapply(candidates,
    function(h) {
    QAW_h <- expit(logit(QAW) + h * (2 * A - 1) / lGAW)
    -mean(Y * log(QAW_h) + (1 - Y) * log(1 - QAW_h))
        })
idx_min <- which.min(risk)
idx_zero <- which.min(abs(candidates)) [1]
labels <- c(expression(R[n] (bar(Q)[list(n,hn)]^list(o))),
        expression(R[n](bar(Q)[list(n,0)]^list(o))))
risk %>% enframe %>%
    mutate(h = candidates) %>%
    filter(abs(h - h[idx_min]) <= abs(h[idx_min])) %>%
    ggplot() +
    geom_point(aes(x = h, y = value), color = "#CC6666") +
    geom_vline(xintercept = c(0, candidates[idx_min])) +
    geom_hline(yintercept = c(risk[idx_min], risk[idx_zero])) +
    scale_y_continuous(
        bquote("empirical logistic risk, " ~ R[n](bar(Q)[list(n,h)]^list(o))),
        trans = "exp", labels = NULL,
        sec.axis = sec_axis(~ .,
                        breaks = c(risk[idx_min], risk[idx_zero]),
                        labels = labels))
```

Figure 15 reveals how moving away slightly from $h=0$ to the left (i.e., to $h_{n}$ equal to -0.004 , rounded to three decimal places) entails a decrease of the empirical risk. The gain is modest at the scale of the empirical risk, but considerable in terms of inference, as we explain below.

### 10.2.5. Justifying the form of the fluctutation

Let us define $\bar{Q}_{n}^{*} \doteq \bar{Q}_{n, h_{n}}$. We justify in two steps our assertion that moving from $\left.\bar{Q}_{n}\right|_{h=0}=\bar{Q}_{n}$ to $\bar{Q}_{n}^{*}$ along fluctuation $\mathscr{Q}\left(\bar{Q}_{n}, \bar{G}_{n}\right)$ has a considerable impact for the inference of $\psi_{0}$.


FIGURE 15. Representing the evolution of the empirical risk function as $h$ ranges over a grid of values. One sees that the risk at $h=0$ (i.e., the risk of $\bar{Q}_{n}$ ) is larger that the minimal risk, achieved at $h \approx-0.004$ (which must be close to that of the optimal $\left.\bar{Q}_{n, h_{n}}\right)$.

First, we note that there is no need to iterate the updating procedure. Specifically, even if we tried to fluctuate $\bar{Q}_{n}^{*}$ along the fluctuation $\mathscr{Q}\left(\bar{Q}_{n}^{*}, \bar{G}_{n}\right)=\left\{\bar{Q}_{n, h^{\prime}}^{*}: h^{\prime} \in \mathbb{R}\right\}$ defined like $\mathscr{Q}\left(\bar{Q}_{n}, \bar{G}_{n}\right)$ (44) with $\bar{Q}_{n, h_{n}}$ substituted for $\bar{Q}_{n}$, then we would not move at all. This is obvious because there is a one-to-one smooth correspondence between the parameter $h^{\prime}$ indexing $\mathscr{Q}\left(\bar{Q}_{n}^{*}, \bar{G}_{n}\right)$ and the parameter $h$ indexing $\mathscr{Q}\left(\bar{Q}_{n}, \bar{G}_{n}\right)$ (namely, $h^{\prime}=h+h_{n}$ ). Therefore, the derivative of (the real-valued function over $\mathbb{R}$ ) $h \mapsto \mathrm{E}_{P_{n}}\left(L_{y}\left(\bar{Q}_{n, h}\right)(O)\right)$ evaluated at its minimizer $h_{n}$ equals 0 . Equivalently (see there in Appendix C.4.3 for a justification of the last but one equality below),

$$
\begin{align*}
\frac{d}{d h} & \left.\mathrm{E}_{P_{n}}\left(L_{y}\left(\bar{Q}_{n, h}^{*}\right)(O)\right)\right|_{h=0} \\
& =-\left.\frac{1}{n} \frac{d}{d h} \sum_{i=1}^{n}\left(Y_{i} \log \bar{Q}_{n, h}^{*}\left(A_{i}, W_{i}\right)+\left(1-Y_{i}\right) \log \left(1-\bar{Q}_{n, h}^{*}\left(A_{i}, W_{i}\right)\right)\right)\right|_{h=0} \\
& =-\frac{1}{n} \sum_{i=1}^{n}\left(\frac{Y_{i}}{\bar{Q}_{n}^{*}\left(A_{i}, W_{i}\right)}-\frac{1-Y_{i}}{1-\bar{Q}_{n}^{*}\left(A_{i}, W_{i}\right)}\right) \times\left.\frac{d}{d h} \bar{Q}_{n, h}^{*}\left(A_{i}, W_{i}\right)\right|_{h=0} \\
& =-\left.\frac{1}{n} \sum_{i=1}^{n} \frac{Y_{i}-\bar{Q}_{n}^{*}\left(A_{i}, W_{i}\right)}{\bar{Q}_{n}^{*}\left(A_{i}, W_{i}\right) \times\left(1-\bar{Q}_{n}^{*}\left(A_{i}, W_{i}\right)\right)} \frac{d}{d h} \bar{Q}_{n, h}^{*}\left(A_{i}, W_{i}\right)\right|_{h=0} \\
& =-\frac{1}{n} \sum_{i=1}^{n} \frac{2 A_{i}-1}{\ell \bar{G}_{n}\left(A_{i}, W_{i}\right)}\left(Y_{i}-\bar{Q}_{n}^{*}\left(A_{i}, W_{i}\right)\right)=0 . \tag{45}
\end{align*}
$$

Second, let $P_{n}^{*}$ be a law in $\mathscr{M}$ such that the $Q_{W}, \bar{G}$ and $\bar{Q}$ features of $P_{n}^{*}$ equal $Q_{n, W}, \bar{G}_{n}$ and $Q_{n}^{*}$, respectively. In other words, $P_{n}^{*}$ is compatible with $Q_{n, W}, \bar{G}_{n}$ and $\bar{Q}_{n}^{*} .{ }^{25}$ Note that the last equality in (45) equivalently writes as

$$
P_{n} D_{2}^{*}\left(P_{n}^{*}\right)=0 .
$$

[^17]Thus the (direct) fluctuation solves (41). Now, we have already argued in Section 10.1.3 that it also holds that

$$
P_{n} D_{1}^{*}\left(P_{n}^{*}\right)=0 .
$$

Consequently, $P_{n}^{*}$ solves the efficient influence curve equation, that is, satisfies (39). As argued in Section 10.1.2, it thus follows that the plug-in estimator

$$
\psi_{n}^{*} \doteq \Psi\left(P_{n}^{*}\right)=\int\left(\bar{Q}_{n}^{*}(1, w)-\bar{Q}_{n}^{*}(0, w)\right) d Q_{n, W}(w)
$$

is asymptotically linear with influence curve $\mathrm{IC}=D^{*}\left(P_{0}\right)$, under mild conditions. Moreover, by virtue of its plug-in construction, it has the additional property that in finite-samples $\psi_{n}^{*}$ will always obey bounds on the parameter space.

### 10.2.6. Alternative fluctuation

The following exercises will have you consider an alternative means of performing a fluctuation. For $a=0,1$, consider the following fluctuation model for $\bar{Q}^{a} \doteq \bar{Q}(a, \cdot)$ :

$$
\begin{equation*}
\mathscr{Q}^{a}\left(\bar{Q}^{a}\right) \doteq\left\{w \mapsto \bar{Q}_{h}^{a}(w) \doteq \operatorname{expit}(\operatorname{logit}(\bar{Q}(a, w))+h): h \in \mathbb{R}\right\} \subset \mathscr{Q} \tag{46}
\end{equation*}
$$

Let

$$
\mathscr{Q}^{\text {alt }}(\bar{Q}) \doteq\left\{(a, w) \mapsto \bar{Q}_{h_{0}, h_{1}}(a, w) \doteq a \times \bar{Q}_{h_{1}}^{1}(w)+(1-a) \bar{Q}_{h_{0}}^{0}(w): h_{0}, h_{1} \in \mathbb{R}\right\}
$$

be the fluctuation model for $\bar{Q}$ that is implied by the two submodels for $\bar{Q}^{1}$ and $\bar{Q}^{0}$.
Also for a given $\bar{G}$ satisfying that $0<\ell \bar{G}(A, W)<1, P_{0}$-almost surely, and both $a=0,1$, consider the loss function $L_{y, \bar{G}}^{a}$ given by

$$
L_{y, \bar{G}}^{a}(f)(A, W) \doteq \frac{\mathbf{1}\{A=a\}}{\ell \bar{G}(A, W)} L_{y}(f)(O)
$$

for any function $f:\{0,1\} \times[0,1] \rightarrow[0,1]$, where $L_{y}$ is defined in (10.2.4) above. It yields the empirical risk function

$$
\left(h_{0}, h_{1}\right) \mapsto \sum_{a=0,1} \mathrm{E}_{P_{n}}\left(L_{y, \bar{G}}^{a}\left(\bar{Q}_{n, h_{a}}^{a}\right)(O)\right) .
$$

1. Comment on the differences between these fluctuation model and loss function compared to those discussed above.
2. Visualize three elements of $\mathscr{Q}^{\text {alt }}\left(\bar{Q}_{0}\right)$ and $\mathscr{Q}^{\text {alt }}\left(\bar{Q}_{n}\right)$ : choose three values for $\left(h_{0}, h_{1}\right)$ and reproduce Figure 14.
3. Argue that in order to minimize the empirical risk over all $\left(h_{0}, h_{1}\right) \in \mathbb{R}^{2}$, we may minimize over $h_{0} \in \mathbb{R}$ and $h_{1} \in \mathbb{R}$ separately.
4. Visualize how the empirical risk with $\bar{G}=\bar{G}_{n}$ varies for different elements $\bar{Q}_{n, h_{0}, h_{1}}$ of $\mathscr{Q}\left(\bar{Q}_{n}\right)$. For simplicity (and with the justification of problem 3), you may wish to make a separate figure (like Figure 15) for $a=0,1$ that illustrates how the empirical risk varies as a function of $h_{a}$ while setting $h_{1-a}=0$.
5. 2 Justify the validity of $L_{y, \bar{G}}^{a}$ as a loss function for $\bar{Q}_{0}^{a}$ : show that amongst all functions that map $w$ to $(0,1)$, the true risk $\mathrm{E}_{P_{0}}\left(L_{y, \bar{G}}^{a}(f)(O)\right)$ is minimized when $f=\bar{Q}_{0}^{a}$.
6. 2 Argue that your answer to problem 2 also implies that the summed loss function $\sum_{a=0,1} L_{y, \bar{G}}^{a}$ is valid for $\bar{Q}$.
7. 2 Justify the combination of these loss function and fluctuation by repeating the calculation in equation (45) above, mutatis mutandis.

### 10.3. Summary and perspectives

The procedure laid out in Section 10 is called targeted minimum loss-based estimation (TMLE). The nomenclature derives from its logistics. We first generated an (un-targeted) initial estimator of $\Psi\left(P_{0}\right)$ by substituting for $P_{0}$ a law $P_{n}^{\circ}$ compatible with initial estimators of some $\Psi$-specific relevant nuisance parameters. Then through loss-minimization, we built a targeted estimator by substituting for $P_{n}^{\circ}$ a law $P_{n}^{*}$ compatible with the nuisance parameters that we updated in a targeted fashion.

The TMLE procedure was coined in 2006 by Mark van der Laan and Dan Rubin (van der Laan and Rubin, 2006). It has since then been developed and applied in a great variety of contexts. We refer to the monographies (van der Laan and Rose, 2011) and (van der Laan and Rose, 2018) for a rich overview.

In summary, targeted learning bridges the gap between formal inference of finite-dimensional parameters $\Psi\left(P_{0}\right)$ of the law $P_{0}$ of the data, via bootstrapping or influence curves, and dataadaptive, loss-based, machine learning estimation of $\Psi$-specific infinite-dimensional features thereof (the so-called nuisance parameters). A typical example concerns the super learningbased, targeted estimation of effect parameters defined as identifiable versions of causal quantities. The TMLE algorithm integrates completely the estimation of the relevant nuisance parameters by super learning (van der Laan et al., 2007). Under mild assumptions, the targeting step removes the bias of the initial estimators of the targeted effects. The resulting TMLEs enjoy many desirable statistical properties: among others, by being substitution estimators, they lie in the parameter space; they are often double-robust; they lend themselves to the construction of confidence regions.

The scientific community is always engaged in devising and promoting enhanced principles for sounder research through the better design of experimental and nonexperimental studies and the development of more reliable and honest statistical analyses. By focusing on prespecified analytic plans and algorithms that make realistic assumptions in more flexible nonparametric or semiparametric statistical models, targeted learning has been at the forefront of this concerted effort. Under this light, targeted learning notably consists in translating knowledge about the data and underlying data-generating mechanism into a realistic model; in expressing the research question under the form of a statistical estimation problem; in analyzing the statistical estimation problem within the frame of the model; in developing $a d$ hoc algorithms grounded in theory and tailored to the question at stake to map knowledge and data into an answer coupled to an assessment of its trustworthiness.

Quoting (van der Laan and Rose, 2018):


#### Abstract

Over the last decade, targeted learning has been established as a reliable framework for constructing effect estimators and prediction functions. The continued development of targeted learning has led to new solutions for existing problems in many data structures in addition to discoveries in varied applied areas. This has included work in randomized controlled trials, parameters defined by a marginal structural model, case-control studies, collaborative TMLE, missing and censored data, longitudinal data, effect modification, comparative effectiveness research, aging, cancer, occupational exposures, plan payment risk adjustment, and HIV, as well as others. In many cases, these studies compared targeted learning techniques to standard approaches, demonstrating improved performance in simulations and realworld applications.


It is now time to resume our introduction to TMLE and to carry out an empirical investigation of its statistical properties in the context of the estimation of $\psi_{0}$.

### 10.4. Empirical investigation

### 10.4.1. A first numerical application

Let us illustrate the principle of targeted minimum loss estimation by updating the Gcomputation estimator built on the $n=1000$ first observations in obs by relying on $\widehat{\mathscr{A}}_{\bar{Q}, \mathrm{kNN}}$, that is, on the algorithm for the estimation of $\bar{Q}_{0}$ as it is implemented in estimate_Qbar with its argument algorithm set to the built-in kknn_algo (see Section 7.6.2). In Section 9.3, we performed a one-step correction of the same initial estimator.

The algorithm has been trained earlier on this data set and produced the object Qbar_hat_kknn. The following chunk of code prints the initial estimator psin_kknn, its one-step update psin_kknn_os, then applies the targeting step and presents the resulting estimator:

```
(psin_kknn)
#> # A tibble: 1 x 2
#> psi_n sig_n
#> <dbl> <dbl>
#> 1 0.0887 0.00245
(psin_kknn_os)
#> # A tibble: 1 x 3
#> psi_n sig_n crit_n
#> <dbl> <dbl> <dbl>
#> 1 0.0888 0.0162 0.000124
(psin_kknn_tmle <- apply_targeting_step(head(obs, 1e3),
                                    wrapper(Gbar_hat, FALSE),
                                    wrapper(Qbar_hat_kknn, FALSE)))
#> # A tibble: 1 x 3
#> psi_n sig_n crit_n
#> <dbl> <dbl> <dbl>
#> 1 0.0888 0.0162 -1.01e-9
```

In the call to apply_targeting_step we provide (i) the data set at hand (first line), (ii) the estimator Gbar_hat of $\bar{G}_{0}$ that we built earlier by using algorithm $\widehat{\mathscr{A}}_{\bar{G}, 1}$ (second line; see Section 7.4.2), (iii) the estimator Qbar_hat_kknn of $\bar{Q}_{0}$ (third line). Apparently, in this particular example, the one-step correction and targeted correction updater similarly the initial estimator.

### 10.4.2. A computational exploration

1. Consult the man page of function apply_targeting_step (run ?apply_targeting_step) and explain what is the role of its input epsilon.
2. Run the chunk of code below. What does it do? Hint: check out the chunks of code of Sections 7.6.2, 8.2.1 and 10.4.1.
```
epsilon <- seq(-1e-2, 1e-2, length.out = 1e2)
Gbar_hat_w <- wrapper(Gbar_hat, FALSE)
Qbar_kknn <- wrapper(Qbar_hat_kknn, FALSE)
psi_trees_epsilon <- sapply(epsilon, function(h) {
    apply_targeting_step(head(obs, 1e3), Gbar_hat_w,
                                    Qbar_trees, epsilon = h) %>%
        select(psi_n, crit_n) %%% as.matrix
})
idx_trees <- which.min(abs(psi_trees_epsilon[2, ]))
psi_kknn_epsilon <- sapply(epsilon, function(h) {
    apply_targeting_step(head(obs, 1e3), Gbar_hat_w,
                                    Qbar_kknn, epsilon = h) %>%
        select(psi_n, crit_n) %>% as.matrix
})
idx_kknn <- which.min(abs(psi_kknn_epsilon[2, ]))
rbind(t(psi_trees_epsilon), t(psi_kknn_epsilon)) %>% as_tibble %>%
    rename("psi_n" = V1, "crit_n" = V2) %>%
    mutate(type = rep(c("trees", "kknn"), each = length(epsilon))) %>%
    ggplot() +
    geom_point(aes(x = crit_n, y = psi_n, color = type)) +
    geom_vline(xintercept = 0) +
    geom_hline(yintercept = c(psi_trees_epsilon[1,idx_trees],
                                    psi_kknn_epsilon[1,idx_kknn])) +
    labs(x = bquote(paste(P[n] ~D~"*", (P[list(n,h)]^o))),
            y = bquote(Psi(P[list(n,h)]^o)))
```

3. Discuss the fact that the colored curves in Figure 16 look like segments.

### 10.4.3. Empirical investigation

To assess more broadly what is the impact of the targeting step, let us apply it to the estimators that we built in Section 8.4.3. In Section 9.3, we applied the one-step correction to the same estimators.

The object learned_features_fixed_sample_size already contains the estimated features of $P_{0}$ that are needed to perform the targeting step on the estimators $\psi_{n}^{d}$ and $\psi_{n}^{e}$, thus we merely have to call the function apply_targeting_step.


FIGURE 16. Figure produced when running the chunk of code from problem 2 in Section 10.4.2.

```
psi_tmle <- learned_features_fixed_sample_size %>%
    mutate(tmle_d =
        pmap(list(obs, Gbar_hat, Qbar_hat_d),
                        ~ apply_targeting_step(as.matrix(..1),
                        wrapper(..2, FALSE),
                                wrapper(..3, FALSE))),
            tmle_e =
                        pmap(list(obs, Gbar_hat, Qbar_hat_e),
                        ~ apply_targeting_step(as.matrix(..1),
                                    wrapper(..2, FALSE),
                                    wrapper(..3, FALSE)))) %>%
    select(tmle_d, tmle_e) %>%
    pivot_longer(c('tmle_d', 'tmle_e'),
                            names_to = "type", values_to = "estimates") %>%
    extract(type, "type", "_([de])$") %>%
    mutate(type = pasteO(type, "_targeted")) %>%
    unnest(estimates) %>%
    group_by(type) %>%
    mutate(sig_alt = sd(psi_n)) %>%
    mutate(clt_ = (psi_n - psi_zero) / sig_n,
            clt_alt = (psi_n - psi_zero) / sig_alt) %>%
    pivot_longer(c('clt_', 'clt_alt'),
                            names_to = "key", values_to = "clt") %>%
    extract(key, "key", "_(.*)$") %>%
    mutate(key = ifelse(key == "", TRUE, FALSE)) %>%
    rename("auto_renormalization" = key)
(bias_tmle <- psi_tmle %>%
    group_by(type, auto_renormalization) %>%
    summarize(bias = mean(clt)) %>% ungroup)
```

```
#> # A tibble: 4 x 3
#> type auto_renormalization bias
#> <chr> <lgl> <dbl>
#> 1 d_targeted FALSE -0.00706
#> 2 d_targeted TRUE -0.0246
#> 3 e_targeted FALSE 0.0353
#> 4 e_targeted TRUE 0.0182
fig <- ggplot() +
    geom_line(aes(x = x, y = y),
        data = tibble(x = seq(-4, 4, length.out = 1e3),
                    y = dnorm(x)),
        linetype = 1, alpha = 0.5) +
    geom_density(aes(clt, fill = type, colour = type),
                psi_tmle, alpha = 0.1) +
    geom_vline(aes(xintercept = bias, colour = type),
            bias_tmle, size = 1.5, alpha = 0.5) +
    facet_wrap(~ auto_renormalization,
            labeller =
                        as_labeller(c('TRUE' = "auto-renormalization: TRUE",
                            'FALSE' = "auto-renormalization: FALSE")),
                        scales = "free")
fig +
    labs(y = "",
        x = bquote(paste(sqrt(n/v[n] {"**")*
                            (psi[n]~{"*"} - psi[0]))))
```

We see that the step of targeting is as promised: the bias of the resulting estimators is minimized relative to the naive estimators. Comparing these results to those obtain using one-step estimators (Section 9.3), we find quite similar performance between one-step and TMLE estimators.

## 11. Closing words

The velocity of advances in machine learning make it an exciting time to work as a statistician. Clearly, statistical inference is more challenging when one considers the sort of infinitedimensional statistical models that underlie these developments. Even defining some fundamental statistical notions, like efficiency, in these settings is a challenge. Constructing estimators that obtain these properties is more challenging still. We hope that this short guide has provided an approachable introduction to this exciting area of research.

Targeted learning is a vibrant and active field of research, with new developments happening along theoretical, applied, and computational axes. The tlverse software environment is actively being developed to provide researchers with new tools for utilizing these methods.


FIGURE 17. Kernel density estimators of the law of two targeted estimators of $\psi_{0}$ (recentered with respect to $\psi_{0}$, and renormalized). The estimators respectively hinge on algorithms $\widehat{\mathscr{A}}_{\bar{Q}, 1}(d)$ and $\widehat{\mathscr{A}}_{\bar{Q}, k N N}(e)$ to estimate $\bar{Q}_{0}$, and on a targeting step. Two renormalization schemes are considered, either based on an estimator of the asymptotic variance (left) or on the empirical variance computed across the iter independent replications of the estimators (right). We emphasize that the $x$-axis ranges differ between the left and right plots.

## Acknowledgements.

David Benkeser and Antoine Chambaz wish to thank a reviewer for her careful reading and encouraging words.

## Appendix A: Notation

$-\doteq$, equal by definition to

- $\mathbf{1}\{S\}$, the indicator of statement $S$, which equals 1 if $S$ is true and 0 otherwise.
- $O \doteq(W, A, Y)$, the generic summary of how one realization of the experiments of interest unfold, our generic observation; $W \in[0,1]$ is the context of action, $A \in\{0,1\}$ is the action undertaken, and $Y \in[0,1]$ is the reward of action $A$ in context $W$. We denote by $\mathscr{O} \doteq$ $[0,1] \times\{0,1\} \times[0,1]$ the set where a generic $O$ takes its values.
- $P, P_{0}, \Pi_{0}, \Pi_{h}, \Pi_{0}^{\prime}, \Pi_{h}^{\prime}$, laws (on $\mathscr{O}$ ) for $O$.
$-P f \doteq \mathrm{E}_{P}(f(O))$ for any law $P$ for $O$ and function $f$ from $\mathscr{O}$ to $\mathbb{R}^{p}$.
- $\|f\|_{P}^{2} \doteq P f^{2}=\mathrm{E}_{P}\left(f(O)^{2}\right)=\int f(o)^{2} d P(o)$, the square of the $L^{2}(P)$-norm of $f$, a function from $\mathscr{O}$ to $\mathbb{R}$.
$-\|f\|_{\infty} \doteq \sup _{o \in \mathscr{O}}|f(o)|$, the essential supremum of $f$, a function from $\mathscr{O}$ to $\mathbb{R}$.
- $P_{n}$, the empirical measure. If the observations are $O_{1}, \ldots, O_{n}$, then $P_{n}$ is a law such that the generic random variable $O$ drawn from $P_{n}$ takes its values in $\left\{O_{1}, \ldots, O_{n}\right\}$ in such a way that $O=O_{i}$ with probability $n^{-1}$ for each $1 \leq i \leq n$.
- $\sqrt{n}\left(P_{n}-P\right)$, where $P_{n}$ is the empirical measure associated to $O_{1}, \ldots, O_{n}$ drawn independently from $P$, the empirical process.
- $X_{n}=o_{P_{0}}(1)$ if $X_{n}$, a random variable built from $O_{1}, \ldots, O_{n}$ independently drawn from $P_{0}$, converges in probability to zero, that is, if $P_{0}\left(\left|X_{n}\right|>t\right)$ converges to zero for all $t>0$ as $n$ goes to infinity. If $n^{c} X_{n}=o_{P_{0}}(1)$, then one also writes $X_{n}=o_{P_{0}}\left(n^{-c}\right)$.
- $X_{n}=O_{P_{0}}(1)$ if $X_{n}$, a random variable built from $O_{1}, \ldots, O_{n}$ independently drawn from $P_{0}$, is bounded in probability, that is if, for all $t>0$ there exists $M>0$ such that $\sup _{n \geq 1} P_{0}\left(\left|X_{n}\right| \geq M\right) \leq t$. If $n^{c} X_{n}=O_{P_{0}}(1)$, then one also writes $X_{n}=O_{P_{0}}\left(n^{-c}\right)$.
- $\mathscr{M}$, the model, that is, the collection of all laws from which $O$ can be drawn and that meet some constraints.
- $\mathscr{M}^{\text {empirical }}$, the collection of all discrete laws on $[0,1] \times\{0,1\} \times[0,1]$, of which $P_{n}$ is a distinguished element.
- $Q_{W}, Q_{0, W}$, marginal laws for $W$ (under $P$ and $P_{0}$, respectively).
$-\bar{G}(W) \doteq \operatorname{Pr}_{P}(A=1 \mid W), \bar{G}_{0}(W) \doteq \operatorname{Pr}_{P_{0}}(A=1 \mid W)$, conditional probabilities of action $A=1$ given $W$ (under $P$ and $P_{0}$, respectively). For each $a \in\{0,1\}, \ell \bar{G}(a, W) \doteq \operatorname{Pr}_{P}(A=a \mid W)$ and $\ell \bar{G}_{0}(a, W) \doteq \operatorname{Pr}_{P_{0}}(A=a \mid W)$.
- $\bar{Q}(A, W)=\mathrm{E}_{P}(Y \mid A, W), \bar{Q}_{0}(A, W)=\mathrm{E}_{P_{0}}(Y \mid A, W)$, the conditional means of $Y$ given $A$ and $W$ (under $P$ and $P_{0}$, respectively).
$-\mathscr{Q} \doteq\{\bar{Q}: P \in \mathscr{M}\}$, the space of regression functions induced by model $\mathscr{M}$.
— $\mathscr{Q}(\bar{Q}, \bar{G}) \subset \mathscr{Q}, \bar{G}$-specific fluctuation model of $\bar{Q}$, see (44).
- $q_{Y}, q_{0, Y}$, conditional densities of $Y$ given $A$ and $W$ (under $P$ and $P_{0}$, respectively).
$-\Psi: \mathscr{M} \rightarrow[0,1]$, given by $\Psi(P) \doteq \int(\bar{Q}(1, w)-\bar{Q}(0, w)) d Q_{W}(w)$, the statistical mapping of interest.
$-\psi \doteq \Psi(P), \psi_{0} \doteq \Psi\left(P_{0}\right)$.
- ${\widehat{\mathscr{A}}, \widehat{\mathscr{A}}_{\bar{G}, 1}}, \widehat{\mathscr{A}}_{\bar{Q}, 1}$, algorithms to be trained on $P_{n}$, i.e., mappings from $\mathscr{M}^{\text {empirical }}$ to the set where lives the feature targeted by the algorithm.
- $\widetilde{\mathscr{A}}_{\bar{G}, s}, \widetilde{\mathscr{A}}_{\bar{Q}, s}, s$-specific oracle algorithms $(s>0)$ that can use the true targeted features $\bar{G}_{0}$ and $\bar{Q}_{0}$ to produce predictions that are almost exact, up to a $N\left(0, s^{2}\right)$ random error term.
- $L_{a}$, the contex-specific logistic (or negative binomial) loss function, given by $-L_{a}(f)(A, W) \doteq A \log f(W)+(1-A) \log (1-f(W))$ for any function $f:[0,1] \rightarrow[0,1]$.
- $L_{y}$, the reward-specific logistic (or negative binomial) loss function, given by $-L_{y}(f)(O) \doteq Y \log f(A, W)+(1-Y) \log (1-f(A, W))$ for any function $f:\{0,1\} \times$ $[0,1] \rightarrow[0,1]$.


## Appendix B: Basic results and their proofs

## B.1. NPSEM

The experiment can also be summarized by a nonparametric system of structural equations: for some deterministic functions $f_{w}, f_{a}, f_{y}$ and independent sources of randomness $U_{w}, U_{a}, U_{y}$,

1. sample the context where the counterfactual rewards will be generated, the action will be undertaken and the actual reward will be obtained, $W=f_{w}\left(U_{w}\right)$;
2. sample the two counterfactual rewards of the two actions that can be undertaken, $Y_{0}=$ $f_{y}\left(0, W, U_{y}\right)$ and $Y_{1}=f_{y}\left(1, W, U_{y}\right) ;$
3. sample which action is carried out in the given context, $A=f_{a}\left(W, U_{a}\right)$;
4. define the corresponding reward, $Y=A Y_{1}+(1-A) Y_{0}$;
5. summarize the course of the experiment with the observation $O=(W, A, Y)$, thus concealing $Y_{0}$ and $Y_{1}$.

## B.2. Identification

Let $\mathbb{P}_{0}$ be an experiment that generates $\mathbb{O} \doteq\left(W, Y_{0}, Y_{1}, A, Y\right)$. We think of $W$ as the context where an action is undertaken, of $Y_{0}$ and $Y_{1}$ as the counterfactual (potential) rewards that actions $a=0$ and $a=1$ would entail, of $A$ as the action carried out, and of $Y$ as the reward received in response to action $A$. Consider the following assumptions:

1. Randomization: under $\mathbb{P}_{0}$, the counterfactual rewards $Y_{0}, Y_{1}$ and action $A$ are conditionally independent given $W$, i.e., $Y_{a} \perp A \mid W$ for $a=0,1$.
2. Consistency: under $\mathbb{P}_{0}$, if action $A$ is undertaken then reward $Y_{A}$ is received, i.e., $Y=Y_{A}$ (or $Y=Y_{a}$ given that $A=a$ ).
3. Positivity: under $\mathbb{P}_{0}$, both actions $a=0$ and $a=1$ have ( $\mathbb{P}_{0}$-almost surely) a positive probability to be undertaken given $W$, i.e., $\operatorname{Pr}_{\mathbb{P}_{0}}\left(\ell \bar{G}_{0}(a, W)>0\right)=1$ for $a=0,1$.

Proposition B.2.1 (Identification). Under the above assumptions, it holds that

$$
\psi_{0}=E_{\mathbb{P}_{0}}\left(Y_{1}-Y_{0}\right)=E_{\mathbb{P}_{0}}\left(Y_{1}\right)-E_{\mathbb{P}_{0}}\left(Y_{0}\right) .
$$

Proof. Set arbitrarily $a \in\{0,1\}$. By the randomization assumption on the one hand (second equality) and by the consistency and positivity assumptions on the other hand (third equality), it holds that

$$
\begin{aligned}
\mathrm{E}_{\mathbb{P}_{0}}\left(Y_{a}\right) & =\int \mathrm{E}_{\mathbb{P}_{0}}\left(Y_{a} \mid W=w\right) d Q_{0, W}(w)=\int \mathrm{E}_{\mathbb{P}_{0}}\left(Y_{a} \mid A=a, W=w\right) d Q_{0, W}(w) \\
& =\int \mathrm{E}_{P_{0}}(Y \mid A=a, W=w) d Q_{0, W}(w)=\int \bar{Q}_{0}(a, W) d Q_{0, W}(w)
\end{aligned}
$$

The stated result easily follows.
Remark. The positivity assumption is needed for $\mathrm{E}_{P_{0}}(Y \mid A=a, W) \doteq \bar{Q}_{0}(a, W)$ to be welldefined.

## B.3. Building a confidence interval

Let $\Phi$ be the standard normal distribution function. Let $X_{1}, \ldots, X_{n}$ be independently drawn from a given law.

## B.3.1. CLT \& Slutsky's lemma

Assume that $\sigma^{2} \doteq \operatorname{Var}\left(X_{1}\right)$ is finite. Let $m \doteq \mathrm{E}\left(X_{1}\right)$ be the mean of $X_{1}$ and $\bar{X}_{n} \doteq n^{-1} \sum_{i=1}^{n} X_{i}$ be the empirical mean. By the central limit theorem (CLT), it holds that $\sqrt{n}\left(\bar{X}_{n}-m\right)$ converges in law as $n$ grows to the centered Gaussian law with variance $\sigma^{2}$.
Moreover, if $\sigma_{n}^{2}$ is a (positive) consistent estimator of $\sigma^{2}$ then, by Slutsky's lemma, $\sqrt{n} / \sigma_{n}\left(\bar{X}_{n}-m\right)$ converges in law to the standard normal law. The empirical variance $n^{-1} \sum_{i=1}^{n}\left(X_{i}-\bar{X}_{n}\right)^{2}$ is such an estimator.
Proposition B.3.1. Under the above assumptions,

$$
\left[\bar{X}_{n} \pm \Phi^{-1}(1-\alpha) \frac{\sigma_{n}}{\sqrt{n}}\right]
$$

is a confidence interval for $m$ with asymptotic level $(1-2 \alpha)$.

## B.3.2. CLT and order statistics

Suppose that the law of $X_{1}$ admits a continuous distribution function $F$. Set $\left.p \in\right] 0,1[$ and, assuming that $n$ is large, find $k \geq 1$ and $l \geq 1$ such that

$$
\frac{k}{n} \approx p-\Phi^{-1}(1-\alpha) \sqrt{\frac{p(1-p)}{n}}
$$

and

$$
\frac{l}{n} \approx p+\Phi^{-1}(1-\alpha) \sqrt{\frac{p(1-p)}{n}} .
$$

Proposition B.3.2. Under the above assumptions, $\left[X_{(k)}, X_{(l)}\right]$ is a confidence interval for $F^{-1}(p)$ with asymptotic level $1-2 \alpha$.

## B.4. Another representation of the parameter of interest

For notational simplicitiy, note that (2a-1) equals 1 if $a=1$ and -1 if $a=0$. Now, for each $a=0,1$,

$$
\begin{aligned}
\mathrm{E}_{P_{0}}\left(\frac{\mathbf{1}\{A=a\} Y}{\ell \bar{G}_{0}(a, W)}\right) & =\mathrm{E}_{P_{0}}\left(\mathrm{E}_{P_{0}}\left(\left.\frac{\mathbf{1}\{A=a\} Y}{\ell \bar{G}_{0}(a, W)} \right\rvert\, A, W\right)\right) \\
& =\mathrm{E}_{P_{0}}\left(\frac{\mathbf{1}\{A=a\}}{\ell \bar{G}_{0}(a, W)} \bar{Q}_{0}(A, W)\right) \\
& =\mathrm{E}_{P_{0}}\left(\frac{\mathbf{1}\{A=a\}}{\ell \bar{G}_{0}(a, W)} \bar{Q}_{0}(a, W)\right)
\end{aligned}
$$

$$
\begin{aligned}
& =\mathrm{E}_{P_{0}}\left(\mathrm{E}_{P_{0}}\left(\left.\frac{1\{A=a\}}{\ell \bar{G}_{0}(a, W)} \bar{Q}_{0}(a, W) \right\rvert\, W\right)\right) \\
& =\mathrm{E}_{P_{0}}\left(\left.\frac{\ell \bar{G}_{0}(a, W)}{\ell \bar{G}_{0}(a, W)} \bar{Q}_{0}(a, W) \right\rvert\, W\right) \\
& =\mathrm{E}_{P_{0}}\left(\bar{Q}_{0}(a, W)\right)
\end{aligned}
$$

where the first, fourth and sixth equalities follow from the tower rule ${ }^{26}$, and the second and fifth hold by definition of the conditional expectation. This completes the proof.

## B.5. The delta-method

Let $f$ be a map from $\Theta \subset \mathbb{R}^{p}$ to $\mathbb{R}^{q}$ that is differentiable at $\theta \in \Theta$. Let $X_{n}$ be a random vector taking its values in $\Theta$.

Proposition B.5.1. If $\sqrt{n}\left(X_{n}-\theta\right)$ converges in law to the Gaussian law with mean $\mu$ and covariance matrix $\Sigma$, then $\sqrt{n}\left(f\left(X_{n}\right)-f(\theta)\right)$ converge in law to the Gaussian law with mean $\nabla f(\theta) \times \mu$ and covariance matrix $\nabla f(\theta) \times \Sigma \times \nabla f(\theta)^{\top}$. In addition, if $\Sigma_{n}$ estimates $\Sigma$ consistently then, by Slutsky's lemma, the asymptotic variance of $\sqrt{n}\left(f\left(X_{n}\right)-f(\theta)\right)$ is consistently estimated with $\nabla f\left(X_{n}\right) \times \Sigma_{n} \times \nabla f\left(X_{n}\right)^{\top}$.

## B.6. The oracle logistic risk

First, let us recall the definition of the Kullback-Leibler divergence between Bernoulli laws of parameters $p, q \in] 0,1[$ :

$$
\mathrm{KL}(p, q) \doteq p \log \left(\frac{p}{q}\right)+(1-p) \log \left(\frac{1-p}{1-q}\right)
$$

It satisfies $\operatorname{KL}(p, q) \geq 0$ where the equality holds if and only if $p=q$.
Let $f:[0,1] \times\{0,1\} \times[0,1] \rightarrow[0,1]$ be a (measurable) function. Applying the tower rule shows that the oracle logistic risk satisfies

$$
\begin{align*}
\mathrm{E}_{P_{0}}\left(L_{y}(f)(O)\right) & =\mathrm{E}_{P_{0}}\left(-\bar{Q}_{0}(A, W) \log f(A, W)-\left(1-\bar{Q}_{0}(A, W)\right) \log (1-f(A, W))\right) \\
& =\mathrm{E}_{P_{0}}\left(\operatorname{KL}\left(\bar{Q}_{0}(A, W), f(A, W)\right)\right)+\text { constant } \tag{47}
\end{align*}
$$

where the above constant equals

$$
-\mathrm{E}_{P_{0}}\left(\bar{Q}_{0}(A, W) \log \bar{Q}_{0}(A, W)-\left(1-\bar{Q}_{0}(A, W)\right) \log \left(1-\bar{Q}_{0, W}(A, W)\right)\right)
$$

In light of (47), $\bar{Q}_{0}$ minimizes $f \mapsto \mathrm{E}_{P_{0}}\left(L_{y}(f)(O)\right)$ over the set of (measurable) functions mapping $[0,1] \times\{0,1\} \times[0,1]$ to $[0,1]$. Moreover, as an average of measures of discrepancy, $\mathrm{E}_{P_{0}}\left(L_{y}(f)(O)\right)$ is also a measure of discrepancy.

[^18]
## Appendix C: More results and their proofs

## C.1. Estimation of the asymptotic variance of an estimator

## C.1.1. IPTW estimator based on a well-specified model

Sketch of proof. The IPTW estimator $\psi_{n}^{b}$ relies on algorithm $\widehat{\mathscr{A}}_{\bar{G}, 1}$, which is "well-specified" in the sense that its output $\bar{G}_{n} \doteq \widehat{\mathscr{A}}_{\bar{G}, 1}\left(P_{n}\right)$ minimizes the empirical risk over a finite-dimensional, identifiable, well-specified working model for $\bar{G}_{0}$. If one introduces $D$ given by

$$
D(O) \doteq \frac{(2 A-1)}{\ell \bar{G}_{0}(A, W)} Y
$$

then the influence curve of $\psi_{n}^{b}$ equals $D-\Psi\left(P_{0}\right)$ minus the projection of $D$ onto the tangent space of the above parametric model for $\bar{G}_{0}$. The variance of the influence curve is thus smaller than that of $D$, hence the conservativeness.

## C.1.2. G-computation estimator based on a well-specified model

Sketch of proof (see (van der Laan and Rose, 2011) page 527). Consider a G-computation estimator $\psi_{n}$ that relies on an algorithm $\widehat{\mathscr{A}}_{\bar{Q}}$ that is "well-specified" in the sense that its output $\bar{Q}_{n} \doteq \widehat{\mathscr{A}}_{\bar{Q}}\left(P_{n}\right)$ minimizes the empirical risk over a finite-dimensional, identifiable, well-specified working model for $\bar{Q}_{0}$. If one introduces $D$ given by

$$
D(O) \doteq \bar{Q}_{0}(1, W)-\bar{Q}_{0}(0, W)
$$

then the influence curve of $\psi_{n}$ equals $D-\Psi\left(P_{0}\right)$ plus a function of $O$ that is orthogonal to $D-\Psi\left(P_{0}\right)$. Thus the variance of the influence curve is larger than that of $D$, hence the anticonservativeness.

## C.2. 2 General analysis of plug-in estimators

Recall that $\widehat{\mathscr{A}}_{Q_{W}}$ is an algorithm designed for the estimation of $Q_{0, W}$ (see Section 7.3) and that we denote by $Q_{n, W} \doteq \widehat{\mathscr{A}}_{Q_{W}}\left(P_{n}\right)$ the output of the algorithm trained on $P_{n}$. Likewise, $\widehat{\mathscr{A}_{\bar{G}}}$ and $\widehat{\mathscr{A}_{\bar{Q}}}$ are two generic algorithms designed for the estimation of $\bar{G}_{0}$ and of $\bar{Q}_{0}$ (see Sections 7.4 and 7.6), $\bar{G}_{n} \doteq \widehat{\mathscr{A}_{\bar{G}}}\left(P_{n}\right)$ and $\bar{Q}_{n} \doteq \widehat{\mathscr{A}_{\bar{Q}}}\left(P_{n}\right)$ are their outputs once trained on $P_{n}$.

Let us now introduce $P_{n}^{\circ}$ a law in $\mathscr{M}$ such that the $Q_{W}, \bar{G}$ and $\bar{Q}$ features of $P_{n}^{\circ}$ equal $Q_{n, W}, \bar{G}_{n}$ and $\bar{Q}_{n}$, respectively. We say that any such law is compatible with $Q_{n, W}, \bar{G}_{n}$ and $\bar{Q}_{n}$.

## C.2.1. Main analysis

Substituting $P_{n}^{\circ}$ for $P$ in (19) yields (34):

$$
\begin{equation*}
\sqrt{n}\left(\Psi\left(P_{n}^{\circ}\right)-\Psi\left(P_{0}\right)\right)=-\sqrt{n} P_{0} D^{*}\left(P_{n}^{\circ}\right)+\sqrt{n} \operatorname{Rem}_{P_{0}}\left(P_{n}^{\circ}\right), \tag{48}
\end{equation*}
$$

an equality that we rewrite as

$$
\begin{equation*}
\sqrt{n}\left(\Psi\left(P_{n}^{\circ}\right)-\Psi\left(P_{0}\right)\right)=-\sqrt{n} P_{n} D^{*}\left(P_{n}^{\circ}\right)+\sqrt{n}\left(P_{n}-P_{0}\right) D^{*}\left(P_{0}\right) \tag{49}
\end{equation*}
$$

$$
\begin{equation*}
+\sqrt{n}\left(P_{n}-P_{0}\right)\left[D^{*}\left(P_{n}^{\circ}\right)-D^{*}\left(P_{0}\right)\right]+\sqrt{n} \operatorname{Rem}_{P_{0}}\left(P_{n}^{\circ}\right) \tag{50}
\end{equation*}
$$

Let us know study in turn the four terms in the above right-hand side sum. Recall that $X_{n}=$ $o_{P_{0}}(1)$ means that $P_{0}\left(\left|X_{n}\right|>t\right)$ converges to zero for all $t>0$ as $n$ goes to infinity.

1. In view of (22), the fourth term is $o_{P_{0}}(1)$ provided that $\sqrt{n}\left\|\bar{Q}-\bar{Q}_{0}\right\|_{P_{0}} \times \|(\bar{G}-$ $\left.\bar{G}_{0}\right) / \ell \bar{G}_{0} \|_{P_{0}}=o_{P_{0}}(1)$. This is the case if, for instance, $\ell \bar{G}_{0}$ is bounded away from zero, and both $n^{1 / 4}\left\|\bar{Q}-\bar{Q}_{0}\right\|_{P_{0}}$ and $n^{1 / 4}\left\|\bar{G}-\bar{G}_{0}\right\|_{P_{0}}$ are $o_{P_{0}}(1)$. What really matters, remarkably, is the product of the two norms. If each norm goes to zero at rate $n^{1 / 4}$, then their product does at rate $\sqrt{n}$. Of course, if one goes to zero at rate $n^{1 / 4+c}$ for some $0<c<1 / 4$, then it suffices that the other go to zero at rate $n^{1 / 4-c}$. See also Section C.3.
2. A fundamental result from empirical processes theory gives us conditions guaranteeing that the third term is $o_{P_{0}}(1)$. By Lemma 19.24 in (van der Vaart, 1998), this is the case indeed if $\left\|D^{*}\left(P_{n}^{\circ}\right)-D^{*}\left(P_{0}\right)\right\|_{P_{0}}=o_{P_{0}}(1)$ (that is, if $D^{*}\left(P_{n}^{\circ}\right)$ estimates consistently $\left.D^{*}\left(P_{0}\right)\right)$ and if $D^{*}\left(P_{n}^{\circ}\right)$ falls (with probability tending to one) into a Donsker class (meaning that the random $D^{*}\left(P_{n}^{\circ}\right)$ must belong eventually to a set that is not too large). Requesting that $\left\|D^{*}\left(P_{n}^{\circ}\right)-D^{*}\left(P_{0}\right)\right\|_{P_{0}}=o_{P_{0}}(1)$ is not much if one is already willing to assume that $n^{1 / 4}\left\|\bar{Q}-\bar{Q}_{0}\right\|_{P_{0}}$ and $n^{1 / 4}\left\|\bar{G}-\bar{G}_{0}\right\|_{P_{0}}$ are $o_{P_{0}}(1)$. Moreover, the second condition can be interpreted as a condition on the complexity/versatility of algorithms $\widehat{\mathscr{A}}_{\bar{G}}$ and $\widehat{\mathscr{A}} \overline{\bar{Q}}$.
3. By the central limit theorem, the second term converges in law to the centered Gaussian law with variance $P_{0} D^{*}\left(P_{0}\right)^{2}$.
4. As for the first term, all we can say is that it is a potentially large (because of the $\sqrt{n}$ renormalization factor) bias term.

## C.2.2. Estimation of the asymptotic variance

Let us show now that, under the assumptions we made in Section C.2.1 and additional assumptions of similar nature, $P_{n} D^{*}\left(P_{n}^{\circ}\right)^{2}$ estimates consistently the asymptotic variance $P_{0} D^{*}\left(P_{0}\right)^{2}$. The proof hinges again on a decomposition of the difference between the two quantities as a sum of three terms:

$$
\begin{align*}
P_{n} D^{*}\left(P_{n}^{\circ}\right)^{2}-P_{0} D^{*}\left(P_{0}\right)^{2}= & \left(P_{n}-P_{0}\right)\left(D^{*}\left(P_{n}^{\circ}\right)^{2}-D^{*}\left(P_{0}\right)^{2}\right)  \tag{51}\\
& +\left(P_{n}-P_{0}\right) D^{*}\left(P_{0}\right)^{2}+P_{0}\left(D^{*}\left(P_{n}^{\circ}\right)^{2}-D^{*}\left(P_{0}\right)^{2}\right) . \tag{52}
\end{align*}
$$

We study the three terms in turn. Recall that $X_{n}=o_{P_{0}}(1 / \sqrt{n})$ means that $P_{0}\left(\sqrt{n}\left|X_{n}\right|>t\right)$ converges to zero for all $t>0$ as $n$ goes to infinity.

1. In light of the study of the third term in Section C.2.1, if $\left\|D^{*}\left(P_{n}^{\circ}\right)^{2}-D^{*}\left(P_{0}\right)^{2}\right\|_{P_{0}}=o_{P_{0}}(1)$ and if $D^{*}\left(P_{n}^{\circ}\right)^{2}$ falls (with probability tending to one) into a Donsker class, then the first term is $o_{P_{0}}(1 / \sqrt{n})$. Furthermore, if $D^{*}\left(P_{n}^{\circ}\right)$ falls (with probability tending to one) into a Donsker class, an assumption we made earlier, then so does $D^{*}\left(P_{n}^{\circ}\right)^{2}$. In addition, if $\left\|D^{*}\left(P_{n}^{\circ}\right)-D^{*}\left(P_{0}\right)\right\|_{P_{0}}=o_{P_{0}}(1)$, another assumption we made earlier, and if there exists a constant $c>0$ such that

$$
\begin{equation*}
\sup _{n \geq 1}\left\|D^{*}\left(P_{n}^{\circ}\right)+D^{*}\left(P_{0}\right)\right\|_{\infty} \leq c \tag{53}
\end{equation*}
$$

$P_{0}$-almost surely, then $\left\|D^{*}\left(P_{n}^{\circ}\right)^{2}-D^{*}\left(P_{0}\right)^{2}\right\|_{P_{0}}=o_{P_{0}}(1)$ too because

$$
\begin{equation*}
\left\|D^{*}\left(P_{n}^{\circ}\right)^{2}-D^{*}\left(P_{0}\right)^{2}\right\|_{P_{0}} \leq c\left\|D^{*}\left(P_{n}^{\circ}\right)-D^{*}\left(P_{0}\right)\right\|_{P_{0}} \tag{54}
\end{equation*}
$$

The existence of such a constant $c$ is granted whenever $\ell \bar{G}_{0}$ and $\ell \bar{G}_{n}$ are bounded away from zero. Note that the condition on $\ell \bar{G}_{n}$ can be inforced by us through the specification of algorithm $\widehat{\mathscr{A}_{\bar{G}}}$.
2. By the central limit theorem, $\sqrt{n}$ times the second term converges in law to the centered Gaussian law with variance $\operatorname{Var}_{P_{0}}\left(D^{*}\left(P_{0}\right)(O)^{2}\right)$, which is finite whenever $\ell \bar{G}_{0}$ is bounded away from zero. By Theorem 2.4 in (van der Vaart, 1998), the second term is thus $O_{P_{0}}(1 / \sqrt{n})$ hence $o_{P_{0}}(1)$.
3. Finally, under assumption (53), the absolute value of the third term is smaller than

$$
\begin{equation*}
c P_{0}\left|D^{*}\left(P_{n}^{\circ}\right)-D^{*}\left(P_{0}\right)\right| \leq c\left\|D^{*}\left(P_{n}^{\circ}\right)-D^{*}\left(P_{0}\right)\right\|_{P_{0}}=o_{P_{0}}(1) \tag{55}
\end{equation*}
$$

where the inequality follows from the Cauchy-Schwarz inequality.
In conclusion, $P_{n} D^{*}\left(P_{n}^{\circ}\right)^{2}-P_{0} D^{*}\left(P_{0}\right)^{2}=o_{P_{0}}(1)$, hence the result.

## C.3. Asymptotic negligibility of the remainder term

Recall that $\|f\|_{P}^{2} \doteq \mathrm{E}_{P}\left(f(O)^{2}\right)$ is the $L_{2}(P)$-norm of $f$, a measurable function from $\mathscr{O}$ to $\mathbb{R}$. Assume that for $a=0,1, \ell \bar{G}_{n}(a, W) \geq \delta>0 Q_{0, W}$-almost everywhere.

The Cauchy-Schwarz inequality then implies that, for $a=0,1$,

$$
\operatorname{Rem}_{P_{0}}\left(P_{n}^{\circ}\right) \leq \frac{2}{\delta} \max _{a=0,1}\left(\left\|\bar{Q}_{n}(a, \cdot)-\bar{Q}_{0}(a, \cdot)\right\|_{P_{0}}\right) \times\left\|\bar{G}_{n}-\bar{G}_{0}\right\|_{P_{0}}
$$

Therefore, if for $a=0,1$,

$$
\left\|\bar{Q}_{n}(a, \cdot)-\bar{Q}_{0}(a, \cdot)\right\|_{P_{0}}=o_{P_{0}}\left(n^{-1 / 4}\right)
$$

and

$$
\left\|\bar{G}_{n}-\bar{G}_{0}\right\|_{P_{0}}=o_{P_{0}}\left(n^{-1 / 4}\right)
$$

then

$$
\operatorname{Rem}_{P_{0}}\left(P_{n}^{\circ}\right)=o_{P_{0}}\left(n^{-1 / 2}\right)
$$

## C.4. Analysis of targeted estimators

## C.4.1. A basic fact on the influence curve equation

Recall the definition of $D_{1}^{*}$ (12). For any estimator $\bar{Q}_{n}^{*}$ of $\bar{Q}_{0}$ and a law $P_{n}^{*}$ that is compatible with $\bar{Q}_{n}^{*}$ and $Q_{n, W}$, it holds that

$$
\begin{aligned}
P_{n} D_{1}^{*}\left(P_{n}^{*}\right) & =\frac{1}{n} \sum_{i=1}^{n} D_{1}^{*}\left(P_{n}^{*}\right)\left(O_{i}\right) \\
& =\frac{1}{n} \sum_{i=1}^{n}\left(\bar{Q}_{n}\left(1, W_{i}\right)-\bar{Q}_{n}\left(0, W_{i}\right)-\int\left(\bar{Q}_{n}(1, w)-\bar{Q}_{n}(0, w)\right) d Q_{n, W}(w)\right) \\
& =\frac{1}{n} \sum_{i=1}^{n}\left(\bar{Q}_{n}\left(1, W_{i}\right)-\bar{Q}_{n}\left(0, W_{i}\right)\right)-\frac{1}{n} \sum_{i=1}^{n}\left(\bar{Q}_{n}\left(1, W_{i}\right)-\bar{Q}_{n}\left(0, W_{i}\right)\right)=0 .
\end{aligned}
$$

## C.4.2. Fluctuation of the regression function along the fluctuation of a law

Let us resume the discussion where we left it at the end of Section 3.3.1. Let $\bar{Q}$ be the conditional mean of $Y$ given $(A, W)$ under $P$. Set arbitrarily $h \in H \backslash\{0\}$ and a measurable function $(w, a) \mapsto f(a, w)$ taking non-negative values. Applying repeatedly the tower rule yields the following equalities:

$$
\begin{aligned}
\mathrm{E}_{P_{h}}(f(A, W) Y)= & \mathrm{E}_{P}(f(A, W) Y(1+h s(O))) \\
= & \mathrm{E}_{P}\left(f(A, W) \mathrm{E}_{P}(Y(1+h s(O)) \mid A, W)\right) \\
= & \mathrm{E}_{P}\left(f(A, W)\left(\bar{Q}(A, W)+h \mathrm{E}_{P}(Y s(O) \mid A, W)\right)\right) \\
= & \mathrm{E}_{P}\left(f(A, W) \frac{\bar{Q}(A, W)+h \mathrm{E}_{P}(Y s(O) \mid A, W)}{1+h \mathrm{E}_{P}(s(O) \mid A, W)}\right. \\
& \left.\quad \times\left(1+h \mathrm{E}_{P}(s(O) \mid A, W)\right)\right) .
\end{aligned}
$$

Now, (9) implies that the density of $(A, W)$ under $P_{h}$ equals $\left(1+h \mathrm{E}_{P}(s(O) \mid A, W)\right)$ when it is evaluated at $(A, W)$. Therefore, the last inequality rewrites as

$$
\mathrm{E}_{P_{h}}(f(A, W) Y)=\mathrm{E}_{P_{h}}\left(f(A, W) \frac{\bar{Q}(A, W)+h \mathrm{E}_{P}(Y s(O) \mid A, W)}{1+h \mathrm{E}_{P}(s(O) \mid A, W)}\right) .
$$

Since this equality is valid for an arbitary $(w, a) \mapsto f(a, w)$ with non-negative values, we can deduce from it that the conditional mean of $Y$ given $(A, W)$ under $P_{h}$ equals

$$
\frac{\bar{Q}(A, W)+h \mathrm{E}_{P}(Y s(O) \mid A, W)}{1+h \mathrm{E}_{P}(s(O) \mid A, W)} .
$$

## C.4.3. Computing the score of a fluctuation of the regression function

Let us resume the discussion where we left it at the beginning of Section 10.2.2. Set $\alpha, \beta \in \mathbb{R}$. The derivative of $h \mapsto \operatorname{expit}(\alpha+\beta h)$ evaluated at $h=0$ satisfies

$$
\left.\frac{d}{d h} \operatorname{expit}(\alpha+\beta h)\right|_{h=0}=\beta \operatorname{expit}(\alpha)(1-\operatorname{expit}(\alpha)) .
$$

Therefore, for any $(w, a) \in[0,1] \times\{0,1\}$,

$$
\begin{aligned}
\left.\frac{d}{d h} \bar{Q}_{h}(a, w)\right|_{h=0} & =\frac{2 a-1}{\ell \bar{G}(a, w)} \operatorname{expit}(\operatorname{logit}(\bar{Q}(a, w)))[1-\operatorname{expit}(\operatorname{logit}(\bar{Q}(a, w)))] \\
& =\frac{2 a-1}{\ell \bar{G}(a, w)} \bar{Q}(a, w)(1-\bar{Q}(a, w))
\end{aligned}
$$

This justifies the last but one equality in (45).
Furthermore the same derivations that led to (45) also imply, mutatis mutandis, that

$$
\left.\frac{d}{d h} L_{y}\left(\bar{Q}_{h}\right)(O)\right|_{h=0}=\frac{2 A-1}{\ell \bar{G}(A, W)}(Y-\bar{Q}(A, W)) .
$$

In this light, and in view of (10), we can think of $\mathscr{Q}(\bar{Q}, \bar{G})$ as a fluctuation of $\bar{Q}$ in the direction of

$$
(w, a, y) \mapsto \frac{2 a-1}{\ell \bar{G}(a, w)}(y-\bar{Q}(a, w))
$$

Thus if $P \in \mathscr{M}$ is such that $\mathrm{E}_{P}(Y \mid A, W)=\bar{Q}(A, W)$ and $P(A=1 \mid W)=\bar{G}(W)$, then we can also think of $\mathscr{Q}(\bar{Q}, \bar{G})$ as a fluctuation of $\bar{Q}$ in the direction of the second component $D_{2}^{*}(P)$ of the efficient influence curve $D^{*}(P)$ of $\Psi$ at $P$ (12).

## References

Barrett, M. (2018). ggdag: Analyze and Create Elegant Directed Acyclic Graphs. R package version 0.1.0.
Benkeser, D. and Chambaz, A. (2020). Causalité et statistique, chapter Une brève introduction à l'apprentissage ciblé, pages 69-119. Éditions Technip. Edited by F. Bertrand, G. Saporta, and C. Thomas-Agnan.
Kuhn, M. (2020). caret: Classification and Regression Training. R package version 6.0-85.
Le Cam, L. M. (1969). Théorie asymptotique de la décision statistique. Séminaire de Mathématiques Supérieures, No. 33 (Été, 1968). Les Presses de l'Université de Montréal, Montreal, Que.
Pfanzagl, J. (1982). Contributions to a general asymptotic statistical theory, volume 13 of Lecture Notes in Statistics. Springer-Verlag, New York-Berlin.
van der Laan, M. J., Polley, E. C., and Hubbard, A. E. (2007). Super learner. Stat. Appl. Genet. Mol. Biol., 6:Art. 25, 23.
van der Laan, M. J. and Rose, S. (2011). Targeted learning. Springer Series in Statistics. Springer, New York. Causal inference for observational and experimental data.
van der Laan, M. J. and Rose, S. (2018). Targeted learning in data science. Springer Series in Statistics. Springer, Cham. Causal inference for complex longitudinal studies.
van der Laan, M. J. and Rubin, D. (2006). Targeted maximum likelihood learning. Int. J. Biostat., 2:Art. 11, 40.
van der Vaart, A. W. (1998). Asymptotic statistics, volume 3 of Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge.
Wickham, H. and Grolemund, G. (2016). R for data science: import, tidy, transform, visualize, and model data. O'Reilly Media, Inc.


[^0]:    1 Department of Biostatistics and Bioinformatics, Emory University, Rollins School of Public Health, 1518 Clifton Road, Atlanta, Georgia 30322.
    E-mail: benkeser@emory. edu
    2 MAP5 (UMR CNRS 8145), Université de Paris, 45 rue des Saints-Pères, 75270 Paris cedex 06. E-mail: antoine.chambaz@u-paris.fr
    ${ }^{1}$ See (Benkeser and Chambaz, 2020) for a shorter presentation of this project, in French.

[^1]:    ${ }^{2}$ A summary of the notation used throughout the text is presented there, in Appendix A.

[^2]:    3 We fine-tuned the marginal law $Q_{0, W}$ of $W$ to make it easier later on to drive home important messages.

[^3]:    ${ }^{4}$ The empirical measure $P_{n}$ is the law such that (i) $X$ drawn from $P_{n}$ takes its values in $\left\{X_{1}, \ldots, X_{n}\right\}$, and (ii) $X=X_{i}$ with probability $n^{-1}$
    5 There are many interesting parameters $\Theta$ for which $\Theta\left(P_{n}\right)$ is not defined, see for instance (6), our parameter of main interest.

[^4]:    6 That is, $P_{h}$ is dominated by $P$ : if an event $A$ satisfies $P(A)=0$, then necessarily $P_{h}(A)=0$ too. Because $P_{h} \ll P$, the law $P_{h}$ has a density with respect to $P$, meaning that there exists a (measurable) function $f$ such that $P_{h}(A)=$ $\int_{o \in A} f(o) d P(o)$ for any event $A$. The function is often denoted $d P_{h} / d P$.

[^5]:    7 Interestingly, if a fluctuation $\left\{P_{h}: h \in H\right\}$ satisfies (10) for a direction $s$ such that $s \neq 0, \mathrm{E}_{P}(s(O))=0$ and $\operatorname{Var}_{P}(s(O))<\infty$, then $h \mapsto \Psi\left(P_{h}\right)$ is still differentiable at $h=0$ with a derivative equal to (11) beyond fluctuations of the form (9).
    8 This may be at first surprising given the parallel drawn in Section 3.3.3 to Euclidean geometry. However, it is important to remember that the model dictates fluctuations of $P$ that are valid submodels with respect to the full model. In turn, this determines the possible directions from which we may approach $P$. Thus, depending on the direction, (11) may hold with different choices of $D^{*}$.

[^6]:    ${ }^{9}$ We can view $\psi_{n}$ as the by product of an algorithm $\widehat{\Psi}$ trained on independent observations $O_{1}, \ldots, O_{n}$ drawn from $P$. We say that the estimator is regular at $P$ if, for any direction $s \neq 0$ such that $\mathrm{E}_{P}(s(O))=0$ and $\operatorname{Var}_{P}(s(O))<\infty$ and fluctuation $\left\{P_{h}: h \in H\right\}$ satisfying (10), the estimator $\psi_{n, 1 / \sqrt{n}}$ of $\Psi\left(P_{1 / \sqrt{n}}\right)$ obtained by training $\widehat{\Psi}$ on independent observations $O_{1}, \ldots, O_{n}$ drawn from $P_{1 / \sqrt{n}}$ is such that $\sqrt{n}\left(\psi_{n, 1 / \sqrt{n}}-\Psi\left(P_{1 / \sqrt{n}}\right)\right)$ converges in law to a limit that does not depend on $s$.

[^7]:    ${ }^{10}$ This also trivially follows from (21).

[^8]:    ${ }^{11}$ We could have used the alternative expression IPAW, where A (like action) is substituted for T (like treatment).

[^9]:    ${ }^{12}$ In fact, if one knows nothing about the feature, then it is certain that it does not belong to whichever small finitedimensional working model we may come up with.
    ${ }^{13}$ See also the companion function compute_1Gbar_hat_AW (run ?compute_lGbar_hat_AW to see its man page.

[^10]:    ${ }^{14}$ Specifically, argument kmax (maximum number of neighbors considered) is set to 5, argument distance (parameter of the Minkowski distance) is set to 2, and argument kernel is set to gaussian. The best value of $k$ is chosen between 1 and kmax by leave-one-out. No outer cross-validation is needed.

[^11]:    ${ }^{15}$ Specifically, argument mstop (number of boosting iterations for prediction) is one among 10, 20 and 30; argument nu (stepsize of the shrinkage parameter) is one among 0.1 and 0.2 ; argument maxdepth (maximum depth of the base learner, a tree) is one among 1,2 and 5 . An outer 10 -fold cross-validation is carried out to select the best combination of fine-tune parameters.

[^12]:    ${ }^{18}$ That is, replicate the construction of $\psi_{n}^{c}$ many times based on data sets obtained by resampling from the original data set, then estimate the asymptotic variance with the empirical variance of $\psi_{n}^{c}$ computed across the replications.

[^13]:    ${ }^{19}$ In words, $v_{n}^{d}$ converges to a lower-bound of the true asymptotic variance.

[^14]:    ${ }^{20}$ We use the expression "tend to" because controlling for multiple testing makes it impossible to make a firm statement.

[^15]:    $\overline{{ }^{21}}$ Note that the algorithm does not really to be trained.

[^16]:    22 or (42).
    ${ }^{23}$ or $P_{n} D^{*}\left(P_{n}^{*}\right)=o_{P_{0}}(1 / \sqrt{n})$.
    ${ }^{24}$ In fact, more generally, the conditional law of $Y$ given $(A, W)$ is not fluctuated. See the corresponding problem in Section 10.2.3.

[^17]:    ${ }^{25}$ In Section 9.1 , we defined $P_{n}^{\circ}$ similarly with $\bar{Q}_{n}$ substituted for $\bar{Q}_{n}^{*}$.

[^18]:    ${ }^{26}$ For any random variable $(U, V)$ such that $\mathrm{E}(U \mid V)$ and $\mathrm{E}(U)$ are well defined, it holds that $\mathrm{E}(\mathrm{E}(U \mid V))=\mathrm{E}(U)$.

