



AgEcon SEARCH

RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.



Estimating and evaluating personalized treatment recommendations from randomized trials with ptr

Matthias Pierce
Division of Psychology and Mental Health
University of Manchester
Manchester, UK
matthias.pierce@manchester.ac.uk

Richard Emsley
Institute of Psychiatry, Psychology and Neuroscience
King's College London
London, UK
richard.emsley@kcl.ac.uk

Abstract. One of the targets of personalized medicine is to provide treatment recommendations using patient characteristics. We present the command `ptr`, which both predicts a personalized treatment recommendation algorithm and evaluates its effectiveness versus an alternative regime, using randomized trial data. The command allows for multiple (continuous or categorical) biomarkers and a binary or continuous outcome. Confidence intervals for the evaluation parameter are provided using bootstrap resampling.

Keywords: `st0639`, `ptr`, personalized treatment recommendations, personalized medicine, stratified medicine

1 Introduction

One of the goals of the modern paradigm of “personalized medicine” is to use data collected on a patient to recommend treatment, rather than treating all patients with the same therapy. A personalized treatment recommendation (PTR) formalizes the process of recommending treatment based on clinically measurable patient traits, termed biomarkers. A PTR is an algorithm that maps single or multiple biomarker inputs onto a decision to treat. For example, breast cancer patients with a low Oncotype DX recurrence score, constructed from a combination of 21 genes, were found to be unlikely to benefit from chemotherapy following surgery (Paik et al. 2006).

As with other kinds of prediction models, developing a PTR requires an estimation and an evaluation step. However, estimating a PTR differs from diagnostic or prognostic model development because the object of the inference (whether a patient benefited from treatment) remains strictly unobserved. Several methods have been developed to tackle this prediction problem, and perhaps the most commonly used is to fit a regression

model with a treatment-by-covariate interaction. Once a PTR has been estimated, it is of interest to know whether it improves on a conservative policy, for example, one where everybody is treated. This potential improvement can be estimated prospectively using a new trial, or it can be done retrospectively using a separate dataset from the one used to develop the model.

In this article, we demonstrate the community-contributed `ptr` command to first estimate a PTR using regression modeling and then to separately estimate the benefit of a PTR compared with an alternative policy, using retrospectively collected data. The command is implemented in the context of a two-armed, randomized controlled trial with a single treatment decision, a binary treatment, and continuous or binary outcomes.

2 Notation and preliminaries

Let $A \in \{0, 1\}$ represent a binary treatment variable; \mathbf{X} a vector of biomarkers, measured at a single baseline visit; and Y the outcome. A PTR is an algorithm that takes the baseline biomarkers as input and outputs a recommendation to treatment or control: $\text{PTR}(\mathbf{X}) : \mathbf{X} \rightarrow A \in \{0, 1\}$. Using counterfactual notation (Rubin 1974), we denote the outcome if a subject was treated as $Y^*(1)$ and if he or she was given control as $Y^*(0)$. Following a randomized trial, the observed outcome for the i th subject is

$$Y_i = Y_i^*(1)A_i + Y_i^*(0)(1 - A_i) \quad (1)$$

Analogously, the outcome that would be observed under a given PTR is

$$Y_i^*(\text{PTR}) = Y_i^*(1)\text{PTR}(\mathbf{X}_i) + Y_i^*(0)\{1 - \text{PTR}(\mathbf{X}_i)\} \quad (2)$$

Note that $Y_i^*(\text{PTR})$ is a potential outcome because it is only partially observed. By contrasting (2) with (1), we clearly see that $Y_i^*(\text{PTR})$ is equal to the observed Y when $A_i = \text{PTR}(\mathbf{X}_i)$ (that is, the subject is randomized to the treatment he or she is recommended). We seek a PTR that recommends treatment for values of \mathbf{X} where the expected benefit of treatment is above a given threshold. When higher values of the outcome are considered beneficial, this can be written as

$$\text{PTR}(\mathbf{X}_i) = I[E\{Y^*(1) - Y^*(0) | \mathbf{X}_i\} > \delta] \quad (3)$$

where $I[\cdot]$ is the indicator function and δ is a number for the minimum improvement necessary for treatment to be recommended over the control. In many applications, this value is zero, although it could take another value if there are economic costs or treatment side effects to account for (see VanderWeele et al. [2019] for a summary).

3 Estimating a PTR using regression

It is natural to separate biomarkers into those that are prognostic (\mathbf{X}^{prog}) and predict the treatment-free outcome and those that are effect modifiers (\mathbf{X}^{mod}) and predict

deviations from the average effect of treatment. For a continuous outcome, a model for the expected outcome might therefore be parameterized by

$$E(Y|A, \mathbf{X}) = \alpha_0 + \boldsymbol{\alpha}\mathbf{X}^{\text{prog}} + A(\beta_0 + \boldsymbol{\beta}\mathbf{X}^{\text{mod}}) \quad (4)$$

This can be estimated by fitting a linear regression model that includes main effects for the estimation of the $\boldsymbol{\alpha}$ parameters and biomarker by treatment interactions for estimation of the $\boldsymbol{\beta}$ parameters. In a randomized trial, $E\{Y^*(a)|\mathbf{X}\} = E(Y|A, \mathbf{X})$, and the parameters of (4) can be substituted into (3) to get an estimate of the PTR:

$$\text{PTR}^{\text{reg}}(\mathbf{X}) = I(\beta_0 + \boldsymbol{\beta}\mathbf{X}^{\text{mod}} > \delta) \quad (5)$$

Note that this depends only on the parameters relating to the average treatment effect (β_0) and the effect modifiers ($\boldsymbol{\beta}$). However, for unbiased estimation of the $\boldsymbol{\beta}$ parameters, the main effect of each modifier must be included (that is, all variables in \mathbf{X}^{mod} should also be in \mathbf{X}^{prog}). Also, inclusion of other variables in \mathbf{X}^{prog} will generally improve the efficiency in estimating the $\boldsymbol{\beta}$ parameters.

4 Improvement under a PTR

After one constructs a PTR, it is of interest to evaluate whether there has been an improvement compared with an alternative policy. Natural alternatives might be either a policy where everybody is treated or a policy where everybody is given the control condition. Writing the expected outcome under a policy as $\mu(\cdot)$, we see the improvements compared with treatment and control are, respectively,

$$\theta_T = \mu(\text{PTR})\mu(A = 1)$$

or

$$\theta_C = \mu(\text{PTR})\mu(A = 0)$$

θ_T is the obvious comparator if the active treatment would be recommended in absence of a PTR, and θ_C is the comparator if the control treatment would be preferred. In the following two sections, we describe two estimates of $\boldsymbol{\theta}$: a nonparametric inverse probability weighted (IPW) estimate and the more efficient augmented inverse probability weighted (AIPW) estimate.

4.1 IPW estimate

As noted above, the outcome under a PTR is observed only in subjects who receive the treatment they were recommended: $Y_i^{\text{obs}}(\text{PTR}) = Y_i A_i \text{PTR}(\mathbf{X}_i) + Y_i (1 - A_i) \{1 - \text{PTR}(\mathbf{X}_i)\}$. The outcome remains missing for the remainder. Using standard missing-data theory (Seaman and White 2013), we can make inferences from $\mu^{\text{obs}}(\text{PTR})$ to $\mu(\text{PTR})$ by reweighting the observed sample by the probability that a subject received the treatment he or she was given. This is given by $p(A) = \pi$ if $A = 1$ and $p(A) = 1 - \pi$

if $A = 0$, where π is the propensity score, fixed by design in a randomized trial. This results in the IPW estimate of the outcome under a PTR:

$$\mu^{(\text{PTR})}_{\text{IPW}} = \frac{1}{n} \sum_{i=1}^n \left(\frac{\{\text{PTR}(\mathbf{X}_i)A_i Y_i\}}{\pi} + \frac{[\{1 - \text{PTR}(\mathbf{X}_i)\}(1 - A_i)Y_i]}{(1 - \pi)} \right)$$

Similarly, the outcome under treatment can be estimated as

$$\mu(A = 1)_{\text{IPW}} = \frac{1}{n} \sum_{i=1}^n \frac{(A_i Y_i)}{\pi}$$

and the outcome under control can be estimated as

$$\mu(A = 0)_{\text{IPW}} = \frac{1}{n} \sum_{i=1}^n \frac{(1 - A_i)Y_i}{(1 - \pi)}$$

4.2 AIPW estimate

We can improve the efficiency of the IPW estimate of the outcome under a treatment policy by borrowing information from the regression of baseline parameters on the outcome (Zhang et al. 2012). This results in the AIPW estimate of the outcome under a PTR:

$$\begin{aligned} \mu^{(\text{PTR})}_{\text{AIPW}} &= \mu^{(\text{PTR})}_{\text{IPW}} \\ &- \frac{1}{n} \sum_{i=1}^n \left(\frac{(A_i - \pi)}{\pi} \text{PTR}(\mathbf{X}_i) m(A = 1, \mathbf{X}_i) + \frac{(A_i - \pi)}{(1 - \pi)} \{1 - \text{PTR}(\mathbf{X}_i)\} m(A = 0, \mathbf{X}_i) \right) \end{aligned}$$

This is equal to the IPW estimate minus an augmentation term that improves the asymptotic efficiency. $m(A = a, \mathbf{X}_i)$ is a regression estimate of the expected outcome under treatment or control, conditional on baseline covariates. These models include variables that are predictive of outcome, regardless of treatment assignment. The outcome under treatment is estimated as

$$\mu(A = 1)_{\text{AIPW}} = \mu(A = 1)_{\text{IPW}} - \frac{1}{n} \sum_{i=1}^n \frac{(A_i - \pi)}{\pi} m(A = 1, \mathbf{X}_i)$$

and the outcome under control is estimated as

$$\mu(A = 0)_{\text{AIPW}} = \mu(A = 0)_{\text{IPW}} - \frac{1}{n} \sum_{i=1}^n \frac{(A_i - \pi)}{(1 - \pi)} m(A = 0, \mathbf{X}_i)$$

AIPW estimates are doubly robust, which means that they are unbiased even when the model for $E(Y|A = a, \mathbf{X})$ is misspecified (Bang and Robins 2005).

4.3 Inference for improvement parameter

Inference for θ comes from resampling the data using the bootstrap procedure and reestimating the statistic on each sample (θ^{bs}). The distribution of θ^{bs} can be translated into confidence intervals using one of three methods. The first estimates the standard error using the standard deviation of θ^{bs} and then calculates the confidence intervals by referring this to a normal distribution. The second, the percentile method, calculates the confidence intervals as the 2.5th and 97.5th percentiles of the distribution of θ^{bs} . The final method uses a bias-corrected and accelerated method to account for the fact that the distribution of θ^{bs} may be skewed away from the normal distribution.

4.4 Estimating and evaluating a PTR for a binary outcome

The process for estimating and evaluating a PTR when the outcome is binary is similar to that outlined above, except that in two instances a logistic regression model is used instead of a linear regression model. The first is in the estimation of the parameters used to construct a PTR. The second is when estimating the expected outcome under treatment or control used in the AIPW estimation of θ . In the binary outcome example, the parameter θ is interpreted as the absolute change in event rate under PTR compared with a policy where everybody or nobody is treated.

5 The ptr command

5.1 Syntax

```
ptr outcome treatment [if] [, modifiers(varlist) covariates(varlist)
  delta(#) ptr(varname) contrast(string) bsamples(#) bootstrap_all
  seed(#) augmented outcome_vars(varlist) tx_vars(varlist) test(varname)
  binary higher lower]
```

outcome can be continuous (the default) or binary (specify the option **binary**), and *treatment* is a binary variable coded 0 for the control group and 1 for the treatment group. The command can be used for both estimation and evaluation of a PTR or for the evaluation of a PTR only. If estimating a PTR, then one must specify the **modifiers(varlist)** option. If evaluating only, then one must specify the **ptr(varname)** option. Note that none of the *varlists* used in this command support the **i.** prefix or the use of **#** for interaction variables. Therefore, interaction and categorical indicator variables must be created beforehand.

5.2 Options

`modifiers(varlist)` specifies the modifier variables used as treatment interaction terms in a regression model, used to estimate the PTR. Their main effect will automatically be included.

`covariates(varlist)` specifies the variables to be used as main effects in the regression model that estimates the PTR.

`delta(#)` sets the minimum threshold for improvement to recommend treatment over control. It can take any real value, and the default is `delta(0)`.

`ptr(varname)` identifies the variable name that indicates whether a subject is recommended treatment under a PTR. This is to be coded 1 if the subject is recommended treatment and 0 if the subject is recommended control. If this option is specified, then the command will evaluate the PTR only and ignore any estimation options.

`contrast(string)` specifies the condition to contrast the PTR with. *string* can be either `treat` for the contrast to be with a policy where everybody is treated or `control` for the contrast to be with a policy where everybody is given control. The default is `contrast(treat)`.

`bsamples(#)` specifies the number of bootstrap samples to be used for inference for the theta parameter. The default is `bsamples(1000)`.

`bootstrap_all` specifies that all bootstrap confidence intervals be reported. This includes normal, percentile, and bias corrected and accelerated. The default is to display the normal confidence intervals only. See [R] **bootstrap**.

`seed(#)` specifies the seed number for consistent bootstrap inference. The default is `seed(0)`.

`augmented` specifies that theta be calculated using the efficient and double-robust AIPW estimate.

`outcome_vars(varlist)` specifies the variables to be used in the outcome model part of the AIPW estimate of theta. Not specifying this option will result in the empirical mean being used as input.

`tx_vars(varlist)` specifies the variables to be used in the treatment prediction part of the inverse probability weighted estimate or AIPW estimate of theta. Not specifying this option will result in the empirical mean being used as input.

`test(varname)` specifies the variable that indicates the part of the dataset that is to be used in evaluation of the PTR, indicating that the remainder is to be used in estimating the PTR. The default is to use all the data for both estimation and evaluation of the PTR.

`binary` specifies that the outcome variable be binary. The default is to assume the outcome is continuous.

higher specifies that higher values of the outcome are better. This is the default option for a continuous outcome.

lower specifies that lower values of the outcome are better. This is the default option for a binary outcome.

5.3 Stored results

`ptr` stores the following in `r()`:

Scalars

<code>r(theta)</code>	estimate of θ
<code>r(theta_se)</code>	standard error of θ
<code>r(mu_ptr)</code>	expected outcome under PTR
<code>r(mu_contrast)</code>	expected outcome under contrasted group

`ptr` also creates the variable `_ptr`, which indicates treatment under the PTR, coded 1 if the subject is recommended treatment and 0 if not.

6 Demonstration of command using simulated data

6.1 Continuous outcome

The following demonstrates `ptr` on a simulated dataset with a continuous outcome. First, 300 observations were created, and a binary treatment variable (A) was randomly generated with treatment assignment probability of 0.5. The following variables were generated: one continuous prognostic variable ($P1$, standard normal); two continuous modifier variables ($M1$ and $M2$, standard normal); one three-level categorical variable, signified by two indicator variables ($M3_{l2}$ and $M3_{l3}$, prevalence 20 and 40%, respectively); and an error term (e , standard normal). The dataset was randomly split 50:50 between training and test datasets. A continuous outcome was generated using the following specification:

$$Y = -2 + 5P1 - 4M1 + 1.5M3_{l2} + 0.5M3_{l3} + A(2 - 3M1 - 2M2 + 1.5M3_{l2} + 2M3_{l3}) + e$$

The following syntax estimates and evaluates a PTR using three modifiers and one covariate:

```
. use ptr_demosim
. ptr Y A, modifiers(M1 M2 M3_12 M3_13) covariates(P1) test(test)
ptr command: constructing and evaluating personalised treatment recommendations
Weights for each modifier variable: inference from regression
```

Variable	Weight	Std. Err.	p	[95% Conf. Interval]	
A	1.947738	.2468109	0.000	1.463999	2.431478
M1	-2.945912	.1584282	0.000	-3.256425	-2.635399
M2	-1.93084	.1742249	0.000	-2.272314	-1.589367
M3_12	1.702721	.3390592	0.000	1.038178	2.367263
M3_13	1.603114	.3289466	0.000	.9583921	2.247837

ptr was generated using:

```
I( 1.94774 + -2.94591*M1 + -1.93084*M2 + 1.70272*M3_12 + 1.60311*M3_13 > 0)
```

Evaluating PTR: inference from 1000 bootstrap samples

	Estimate	Std. Err.	p	[95% Conf. Interval]	
mu ptr	-.2042839	.8855018	0.409	-1.939836	1.531268 (N)
mu contrast	.3425118	.8628713	0.346	-1.348685	2.033708 (N)
theta	-.5467956	.7003033	0.217	-1.919365	.8257735 (N)

(N) normal confidence interval

The output shows both estimation and evaluation steps. The first table provides the weights for each variable used in the estimation of the PTR. These are the estimated average treatment effect and the modifier by treatment interaction terms, as per the linear model (4). These are used to generate the PTR algorithm [provided in text below the first table, according to (1) above]. The standard errors, p -values, and confidence intervals of the weights are derived from Wald statistics and are the same as those displayed after the `regress` (see [R] `regress`) command.

The second table shows the estimates of the outcome under the PTR (`mu ptr`), the outcome under the contrast scenario (`mu contrast`; everybody receives treatment as the default), and the expected improvement under the PTR (`theta`). These are all estimated on the test dataset, using the default IPW approach. The standard errors, p -values, and confidence intervals are from the bootstrap procedure.

To demonstrate how to evaluate a PTR using the AIPW method, we simulated an arbitrary PTR on the same dataset: $\text{PTR} = I(M1 < 0 \ \& \ M2 < 0)$.

The following syntax evaluates this using the AIPW method, using the covariate $P1$ in the linear model for $m(A, \mathbf{X}_i)$:

```
. generate PTR=M1<0 & M2<0 // new PTR
. ptr Y A, augmented outcome_vars(M1 M2) test(test) ptr(PTR)
When ptr(varname) specified, this program will evaluate the PTR only
ptr command: constructing and evaluating personalised treatment recommendations
Evaluating PTR: inference from 1000 bootstrap samples
```

	Estimate	Std. Err.	p	[95% Conf. Interval]	
mu ptr	.2088446	.8168414	0.399	-1.392135	1.809824 (N)
mu contrast	.9424386	.7287016	0.098	-.4857902	2.370667 (N)
theta	-.733594	.6951825	0.146	-2.096127	.6289387 (N)

(N) normal confidence interval

6.2 Binary outcome

The following demonstration uses the same covariates defined above, except with a binary outcome variable that has probability

$$Y = \text{logit}^{-1} \{-4 - 2P1 - 2M1 + M3_{l2} + A(-3 - 4M1 + 2M2 + 1.5M3_{l2} + 1M3_{l3})\}$$

where logit^{-1} is the inverse logit function. This probability is referred to as a random Bernoulli distribution to simulate the binary outcome.

The estimation and evaluation of the PTR under this distribution is given by the following output. This time, we display all bootstrap standard errors using the option `bootstrap_all` and contrast the outcome under the PTR with the outcome under the control.

```
. ptr Y_bin A, modifiers(M1 M2 M3_12 M3_13) covariates(P1) test(test) binary
> contrast(control) bootstrap_all
ptr command: constructing and evaluating personalised treatment recommendations
Weights for each modifier variable: inference from regression
```

Variable	Weight	Std. Err.	p	[95% Conf. Interval]	
A	-1.060465	1.600099	0.507	-4.196595	2.075664
M1	-3.490736	1.476918	0.018	-6.385436	-.5960362
M2	1.396803	.8581385	0.104	-.2851144	3.07872
M3_12	.9849883	2.180123	0.651	-3.287965	5.257942
M3_13	1.640338	1.591238	0.303	-1.478425	4.7591

ptr was generated using:

```
I( -1.06047 + -3.49074*M1 + 1.39680*M2 + 0.98499*M3_12 + 1.64034*M3_13 < 0)
```

Evaluating PTR: inference from 1000 bootstrap samples

	Estimate	Std. Err.	p	[95% Conf. Interval]	
mu ptr	.1621622	.0432956	0.000	.0773044	.2470199 (N)
				.0839082	.2512303 (P)
				.0876635	.2575454 (BC)
mu contrast	.1486486	.0408638	0.000	.0685572	.2287401 (N)
				.0730637	.2340105 (P)
				.0815545	.2410766 (BC)
theta	-.0135135	.0137815	0.163	-.0405249	.0134978 (N)
				-.0467981	0 (P)
				-.0529471	0 (BC)

(N) normal confidence interval

(P) percentile confidence interval

(BC) bias-corrected confidence interval

7 Discussion

In this article, we explained how a PTR can be estimated and evaluated using data retrospectively collected from two-armed randomized controlled trials. We implemented this theory in a package: `ptr`. This command allows for potentially multiple biomarkers and binary or continuous outcomes. The command also allows users to specify a minimum threshold for improvement. For the evaluation step, the user can specify a doubly robust method that has efficiency gains over the standard inverse probability method. The user can also contrast the outcome under a PTR with several different policies (for example, treat everybody). This is the first available command in Stata; for similar programs in R, see the packages `TreatmentSelection` (Janes et al. 2014) and `PTE` (Kapelner et al. 2014).

Ideally, to avoid overoptimism in the evaluation step, one should use data in the estimation and evaluation steps from two separate randomized controlled trials. There are two further problems with estimating a PTR using a regression model. First, the model may be misspecified in that it omits higher-order, transformed, or interaction terms between the biomarkers that might better explain the data. This can be corrected by specifying a rich model with many parameters. However, this may result in the second problem: overfitting of the model, where the model predicts well in a training dataset

but not well in any other datasets. This may be solved by extending the `ptr` command to incorporate penalized regression, for example, the lasso estimator recently added to Stata 16 (see [LASSO] `lasso`).

We recommend that the `ptr` command be used for exploratory analysis, to understand what form a PTR might take and how well it might perform compared with other regimes. A PTR estimated from retrospectively collected trial data will likely not be fit for direct use in clinical practice. We recommend that a prospective trial should be used to see how well `ptr` performs versus alternative regimes, potentially allowing for the PTR to be adapted as data from the trial are collected (see <http://www.bigted.org/> for details of trial designs to be used for this purpose).

8 Acknowledgments

We acknowledge the guidance and support of our departed mentor Professor Graham Dunn, and we dedicate this to his memory. This work was funded by the MRC North West Hub for Trials Methodology (grant reference: MR/K025635/1).

9 Programs and supplemental materials

To install a snapshot of the corresponding software files as they existed at the time of publication of this article, type

```
. net sj 21-2
. net install st0639      (to install program files, if available)
. net get st0639          (to install ancillary files, if available)
```

10 References

- Bang, H., and J. M. Robins. 2005. Doubly robust estimation in missing data and causal inference models. *Biometrics* 61: 962–973. <https://doi.org/10.1111/j.1541-0420.2005.00377.x>.
- Janes, H., M. D. Brown, Y. Huang, and M. S. Pepe. 2014. An approach to evaluating and comparing biomarkers for patient treatment selection. *International Journal of Biostatistics* 10: 99–121. <https://doi.org/10.1515/ijb-2012-0052>.
- Kapelner, A., J. Bleich, Z. D. Cohen, R. J. DeRubeis, and R. Berk. 2014. Evaluating the effectiveness of personalized medicine with software. ArXiv Working Paper No. arXiv:1404.7844. <http://arxiv.org/abs/1404.7844>.
- Paik, S., G. Tang, S. Shak, C. Kim, J. Baker, W. Kim, M. Cronin, F. L. Baehner, D. Watson, J. Bryant, J. P. Costantino, C. E. Geyer, D. L. Wickerham, and N. Wolmark. 2006. Gene expression and benefit of chemotherapy in women with node-negative, estrogen receptor-positive breast cancer. *Journal of Clinical Oncology* 24: 3726–3734. <https://doi.org/10.1200/JCO.2005.04.7985>.

- Rubin, D. B. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66: 688–701. <https://doi.org/10.1037/h0037350>.
- Seaman, S. R., and I. R. White. 2013. Review of inverse probability weighting for dealing with missing data. *Statistical Methods in Medical Research* 22: 278–295. <https://doi.org/10.1177/0962280210395740>.
- VanderWeele, T. J., A. R. Luedtke, M. J. van der Laan, and R. C. Kessler. 2019. Selecting optimal subgroups for treatment using many covariates. *Epidemiology* 30: 334–341. <https://doi.org/10.1097/EDE.0000000000000991>.
- Zhang, B., A. A. Tsiatis, E. B. Laber, and M. Davidian. 2012. A robust method for estimating optimal treatment regimes. *Biometrics* 68: 1010–1018. <https://doi.org/10.1111/j.1541-0420.2012.01763.x>.

About the authors

Matthias Pierce is a research fellow in the Centre for Women's Mental Health at the University of Manchester.

Richard Emsley is a professor of medical statistics and trials methodology at the Institute of Psychiatry, Psychology and Neuroscience at King's College London.