Double-Blind Randomized Placebo-Controlled Trials in the Treatment of Affective Disorders: Problems and Alternatives

E. Severus, M.D.^{1, *}

E. Laber, Ph.D.²

I. Lipkovich, Ph.D.³

Address

^{*1}Department of Psychiatry and Psychotherapy, University Hospital Carl Gustav Carus, Technische Universität Dresden, Fetscherstraße 74, 01307, Dresden, Germany

Email: Emanuel.Severus@uniklinikum-dresden.de

²NCSU Statistics Department, 2311 Stinson DriveCampus Box 8203, Raleigh, NC 27695-8203, USA

³Quintiles, 4820 Emperor Blvd, 27703, Durham, NC, USA

Published online: 15 July 2015

© Springer International Publishing AG 2015

This article is part of the Topical Collection on *Mood Disorders*

Keywords Marginal structural models · Q-learning · Bipolar disorders · Personalized medicine · Evidence-based medicine

Opinion statement

Double-blind randomized placebo-controlled trials, in which a new approval-seeking drug is compared to placebo, are generally considered to be the gold standard for evaluating the efficacy of a psychopharmacological intervention in the treatment of affective disorders. However, this type of study has substantial limitations regarding the external validity of its results. Therefore, alternatives within randomized controlled trials, such as non-inferiority trials or study designs in which the new drug or placebo is given simultaneously on top of standard antidepressant or anti-manic treatment are of substantial interest. In addition, innovative statistical approaches such as marginal structural models or Q-learning should be employed to evaluate potential causal effects associated with different fixed treatment regimes that are realized in actual clinical practice and help to establish treatment recommendations that are tailored to individual patient characteristics.

Introduction

Double-blind randomized placebo-controlled trials are generally considered to be the gold standard for evaluating the efficacy of a psychopharmacological intervention in the treatment of affective disorders. Furthermore, positive results from a trial of this type are a prerequisite for a new drug to get approval for use in clinical care. Placebo-controlled randomized controlled trials (RCTs) have several strengths. Randomization with allocation concealment makes sure that all study participants have a nonzero chance to be assigned to one of the treatment groups, so that known and unknown risk factors and prognostic factors regarding study outcome should be equally distributed between them. Therefore, if there is a difference in the outcome criteria between the treatment groups at the end of the study, this difference might be interpreted as causally related to differences in the efficacy of the applied interventions in the population of patients represented by the sample of patients in the trial. A placebo is a therapeutic intervention which does not differ from the active intervention regarding mode of application, appearance, color, taste and smell, but which does not have a specific mode of action and is therefore believed not to have an effect on a person. Thus, in randomized placebo-controlled trials, placebo is used as the control condition for the active intervention to get a realistic idea of the factual efficacy of the active intervention compared to an intervention with a substance with no effect. Blinding is used to hide patient assignment to the study arms to participants involved in the study until the end of the study. The goal of this measure is to minimize biases which may evolve if treatment outcome is influenced by the knowledge of treatment assignment (e.g., active drug, placebo). Double-blinding refers to the fact that not only the patient but also the treating physician/rater is not informed about treatment assignment.

All the study procedures described above are measures which aim to ensure that the study has high internal validity, i.e., that the results of the study reflect the true effect of the intervention, free from systematic errors (bias). High internal validity is a requirement for the applicability of the study results in clinical practice. Unfortunately, experience has shown that there are several problems with the interpretation of the results of randomized placebo-controlled trials in the treatment of affective disorders, most of them relate to the external validity of the results of such trials. A study is considered to have high external validity if the study results can be

well transferred into routine clinical care. Differences in patient characteristics, applied interventions or general frame-work may result in a study having low external validity and consequently only limited use for clinical care. For example, patients with a major depressive episode from the Sequenced Treatment Alternatives to Relieve Depression (STAR*D) project fulfilling typical entry criteria for phase III trials (efficacy sample) had a shorter average duration of illness and lower rates of family history of substance abuse, prior suicide attempts, and anxious and atypical symptom features than those patients not fulfilling the inclusion criteria (non-efficacy sample). Furthermore, despite similar medication dosing and time at exit dose, the efficacy sample tolerated the antidepressant used (citalopram) better. In addition, they also had higher rates of response (51.6 versus 39.1 %) and remission (34.4 versus 24.7 %). In addition, these differences persisted even after adjustments for baseline differences [1]. Nowadays, many study participants in phase III trials are recruited by advertisements. There is evidence suggesting that these symptomatic volunteers may not be representative of treatment-seeking patients encountered in clinical care [2].

As mentioned earlier, in the treatment of affective disorders, the knowledge of treatment assignment (active drug, placebo) is believed to have an impact on the outcome of an intervention; therefore, blinding is used to minimize potentially resulting biases. However, knowledge of the possibility of being assigned to placebo, about which the study participants are apprised through the informed consent form (ICF), also influences response and remission rates both for the active antidepressant and the placebo [3, 4] and study outcome [5]. As the use of placebo is not part of routine clinical care, the external evidence of placebo-controlled trials may be compromised. Nevertheless, a greater likelihood of receiving placebo predicts greater antidepressant-placebo separation at endpoint. It has been proposed that informed consent could be even more explicit and clear about the probability of receiving placebo or that study designs could be modified in such a way so as to increase the probability of receiving placebo (versus active treatment) [5]. In addition, such a measure is likely to not attract real treatment-seeking patients representative of those encountered in routine clinical care but to attract symptomatic volunteers with more opaque motives [2].

To address the problems outlined above and to mirror clinical practice in clinical trial design more closely (and thereby improve the external validity of randomized double-blind placebo-controlled trials in the pharmacological treatment of major depression), we have proposed to introduce two key measures as part of a new concept [6]: (i) active medication for all study participants at the end of late phase III clinical trials; (ii) to re-conceptualize the role of the placebo and what constitutes being randomized to the placebo arm. This would involve explaining to study participants in the informed consent form that improvement in their depressive symptoms may occur spontaneously while on the placebo arm, in which case they may have been spared unnecessary medication and to additionally emphasize that study participation is likely to result in increased and more regular contact with clinicians and closer monitoring of their symptoms which can be beneficial in its own right. This approach may also allow to conduct a placebo-controlled trial in major depression in an open fashion [6, 7].

Another type of study that may be useful in avoiding at least some of the problems discussed above is non-inferiority trials which compare a new therapy against standard care. However, in order to demonstrate the non-inferiority (or superiority) of the new therapy versus the established active comparator, a large sample size may be required. However, this type of study does not address the clinically relevant question of whether any treatment is more effective than no treatment, for

example compared to a watchful waiting strategy as used in clinical practice and recommended in international treatment guidelines.

Another interesting strategy is to conduct double-blind placebo-controlled trials in which the new drug or placebo is given simultaneously on top of standard antidepressant or anti-manic treatment. The feasibility of this approach in terms of enrolling real treatment-seeking severely ill patients (some of them in-patients) has been demonstrated recently. In addition, it proved successful in terms of drug-placebo separation at end-point [8]. Such an approach would also allow for the inclusion of treatment-resistant patients or patients with suicidal ideation who are typically excluded from placebo-controlled trials but represent a substantial portion of treatment-seeking patients in routine clinical care.

While some of the approaches outlined above may be useful in addressing typical problems of standard double-blind randomized placebo-controlled trials in the treatment of affective disorders, other problems remain unsolved. These are, among others, how to assess the efficacy of watchful waiting often practiced with mild major depressive episodes in routine clinical care as well as how to personalize treatment to patients with affective disorders.

In the following, we introduce two statistical procedures that will allow us to address, among others, these questions: marginal structural models and Q-learning.

Q-learning

The Q-learning algorithm is based on the point-of-view that the best possible clinical care requires treatment recommendations that are tailored to individual patient characteristics. This notion is formalized as a decision rule that maps patient characteristics to a recommended treatment. An example of a simple treatment regime in the context of bipolar depression [9•] is as follows:

If hypo(manic) prior to depressive episode then anti-depressant alone;
Else if age exceeds 60, then mood stabilizer and paroxetine;
Else mood stabilizer and bupropion.

Thus, a treatment regime is an algorithm that dictates treatment according to patients' current health status. Furthermore, a treatment regime estimated from data can be used to generate new clinical questions for future investigation, e.g., do patients experiencing hypo(mania) prior to a depressive episode benefit from an adjunctive mood stabilizer? Procedures for estimating treatment regimes like that listed in [9•] have been well-studied in statistics. Recent surveys

include Nahum-Shani et al. [10, 11] and Schulte et al. (2014) [12]. Extensions of these methods exist to handle missing data [13], censoring [14, 15], multiple outcomes [16], and continuous treatments (e.g., dose) [17]. However, all of these estimators share the same basic structure. Our goal is to describe the statistical thinking that underpins these estimators and to describe the implementation details of one such estimator, Q-learning.

Suppose that given any regime, say π , you were able to determine the expected clinical outcome if you used π to assign treatments to a population of interest, say $V(\pi)$. Then, to compare two regimes, say π and π' , it would be sufficient to compare $V(\pi)$ and $V(\pi')$, e.g., if the clinical outcome were disease-free survival time, then we would prefer π to π' if $V(\pi) > V(\pi')$. Given any collection of regimes, Π , the optimal regime, in terms of the mean clinical outcome, is the regime that maximizes (or minimizes) $V(\pi)$ over all π in Π . Estimation of an optimal regime from data typically proceeds in three steps: (S1) choose a class of potential regimes, Π ; (S2) derive an estimator $\hat{V}(\pi)$ for each π in Π ; and (S3) choose the maximizer of $\hat{V}(\pi)$ over π in Π as the estimator of the optimal regime [18].¹ The class used in (S1) is maybe infinite, e.g., the class of all possible regimes; however, in some applications, it may be desirable to restrict this class to ensure parsimony, interpretability, or the satisfaction of logistical/cost constraints [17, 19, 20]. Thus, (S1) is generally informed by subject matter knowledge. Implementation of (S2) depends on the study design used to generate the available data; special care must be taken with observational studies due to potential confounding [12]. We describe an implementation of (S2) using Q-learning below. The complexity of (S3) depends on the class chosen in (S1) and the estimator used in (S2). As we show below, (S3) is trivial when using Q-learning with set to be all possible regimes. However, in other settings, this step can be extremely computationally intensive requiring specialized optimization algorithms or heuristics [16, 21].

Mathematical description of Q-learning

We assume the observed data are $\{(X_i, A_i, Y_i)\}_{i=1}^n$, which comprise n independent, identically distributed (*i.i.d.*) copies of the tuple (X, A, Y) where X denotes a vector of pre-treatment patient characteristics, A denotes treatment received coded to take values in $\{0, 1\}$, e.g., active treatment ($A=1$) and control ($A=0$), and Y denotes a scalar outcome coded so that higher values are better. Formally, a treatment regime is a map, π , so that a patient presenting with $X=x$ is assigned treatment $\pi(x)$. To define an optimal treatment regime, we use the language of potential outcomes. Let $Y^*(a)$ denote the potential outcome that a patient would experience if assigned treatment a . For any regime, π , define $Y^*(a) = \pi(X)Y^*(1) + (1 - \pi(X))Y^*(0)$ to be the potential outcome if treatment is assigned according to π . For any π define $V(\pi) = EY^*(\pi)$. The optimal regime, say π^{opt} , satisfies $V(\pi^{\text{opt}}) \geq V(\pi)$ for all regimes π .

To estimate π^{opt} from the data, we make the following assumptions: (A1) the treatment assignment mechanism is completely determined by X and factors that are independent of the potential outcomes $\{Y^*(0), Y^*(1)\}$; (A2)

treatments are randomly assigned such that all patients have probability at least $p_{\min} > 0$ of receiving each treatment; and (A3) the outcome observed is the potential outcome under treatment actually received. We have stated these assumptions informally, precise mathematical versions of these assumptions can be found in Zhang et al. 2012 [21]. Assumptions (A1)–(A2) are satisfied by design in a randomized clinical trial but are not applicable to an observational study.

Define $Q(x, a) = E(Y|X=x, A=a)$. Under (A1)–(A3), it can be shown that for any regime π $V(\pi) = EQ\{X, \pi(X)\}$. From this expression, it follows that $\pi^{\text{opt}}(x) = 1$ if $Q(x, 1) > Q(x, 0)$ and $\pi^{\text{opt}}(x) = 0$ otherwise. Q-learning estimates π^{opt} by first estimating the $Q(x, a)$, say by $\hat{Q}(x, a)$, and subsequently defining the estimator $\hat{\pi}(x) = 1_{\hat{Q}(x, 1) > \hat{Q}(x, 0)}$. We illustrate estimation of $Q(x, a)$ using a linear model fit by least squares. We postulate the linear model $Q(x, a; \beta) = x_0^T \beta_0 + ax_1^T \beta_1$, where x_0 and x_1 summaries constructed from x and $\beta = (\beta_0^T, \beta_1^T)^T$. Q-learning consists of the following two-steps.

(Q1) Compute $\hat{\beta}$ as the minimizer of $\sum_{i=1}^n \{\Upsilon_i - Q(X_i, A_i; \beta)\}^2$.

(Q2) Define $\hat{\pi}(x) = 1_{\hat{Q}(x, 1) > \hat{Q}(x, 0)} = 1_{x_1^T \hat{\beta}_1 > 0}$.

Because Q-learning requires only a least squares fit, it can be implemented readily using essentially any existing statistical computing environment. Under the assumption that $Q(x, a) = Q(x, a; \beta^*)$ for some “true” coefficient vector β^* , confidence intervals and hypothesis tests for the components of β^* can be computed using standard methods for ordinary least squares. Furthermore, in the active-treatment versus control setting, $x_1^T \hat{\beta}_1$ can be interpreted as a score so that only patients with a high-score are identified as candidates for the new treatment.

Marginal structural models

Marginal structural models (MSMs) [22–25] aim at evaluating potential causal effects associated with different fixed treatment regimes that were realized in actual clinical practice by using the framework of potential outcomes and accounting for time-dependent confounding. When comparing outcomes associated with various observed treatment regimes, confounding may occur because a decision to assign patients to different treatments over time may be driven by patients intermediate outcomes observed during the course of the treatment.

For example, a patient who developed a depressive syndrome with sleeping problems precipitated by an acute stressor, may discuss with his/her psychiatrist how to proceed [26•]. Depending on previous experience, they agree that the best option is just to take sleeping medication for a short period, take several days off work, and resume exercising on a regular basis. Suppose that after 1 week, the patient reports better sleep and a modest improvement in their depressive symptoms; however, this improvement does not fulfill the criteria for response. The psychiatrist and the patient therefore decide to continue treatment for another 2 weeks. After these 2 weeks, in spite of the interventions described above, the

depressive symptoms have worsened. After discussing the patient's recent health trajectory, the psychiatrist and the patient agree that it is now time to start with an antidepressant. Two weeks later, the patient realizes a substantial improvement, and after another 2 weeks, the patient fulfills the criteria for response. In this example, the first-line of treatment, denoted *A*, would be "watchful waiting," including temporary use of sleeping medication plus psychosocial interventions (take some days off work, exercising, no formal psychotherapy). The second-line treatment, denoted *B*, would be treatment *A* plus an antidepressant—which was introduced after failure of treatment *A*. The entire treatment sequence over time (also referred to as "treatment plan" or "treatment regime") can be depicted as *AABB* where positions of the letters are associated with consecutive time intervals represented as (start week-end week): (0–1), (1–3), (3–5), and (5–7), beginning from the time of initiation of psychosocial interventions. Associated with this treatment sequence is the sequence of patient's severity scores: Y_0 , Y_1 , Y_3 , Y_5 , and Y_7 , observed at weeks 0 (baseline), 1, 3, 5, and 7. The decisions to assign the patient initially to treatment *A* and to switch them from treatment *A* to *B* after the first 3 weeks of watchful waiting were largely driven by the patients outcomes observed prior to these time points. Therefore, a direct comparison of observed outcomes associated with different treatment plans would be inappropriate as it may lead to biased estimate of relative efficacy of the two plans. For example, imagine another patient with a treatment sequence *BBBB* who may have been very sick from the beginning and therefore put on antidepressant immediately, after which s/he had some modest improvement after 9 weeks. A direct comparison of observed outcomes for plans *AABB* and *BBBB* for these two patients (or, more generally, for a sample of such patients) may be biased in favor of *AABB* and therefore would mask potential efficacy associated with psychopharmacological treatment *B*. In the hypothetical situation of a "careless physician" who would toss a fair coin at each time point to decide whether the patient should remain on their current treatment or switch to the other one, a simple direct comparison of observed outcomes for each realized treatment sequence would be an unbiased estimate of associated treatment effect by the virtue of randomization. Now, imagine that our hypothetical physician tosses a biased coin so that the probability of treatment assignment (*A* or *B*) depends on the severity of a patient's observed depressive symptoms so that more severe patients had a larger probability of being assigned to treatment *B*. Specifically, let us assume that a patient who is current receiving treatment *A* has a probability 0.9 of being switched from *A* to *B* if their depression severity score exceeds a relapse cutoff, c_1 , and has a probability 0.3 of switching to *B* if their depression score is below c_1 , whereas a patient who is already receiving treatment *B* would have a probability of 0.7 to remain on *B* if their severity score is above a remission cutoff, c_2 , and a probability of 0.5 of remaining on *B* if it is below c_2 . A similar treatment assignment mechanism was used in a simulation experiment described in Severus et al. 2013 [26•]. Note that had we known these probabilities we could have utilized them to estimate for each patient the probability of being assigned to the treatment sequence that they actually received during the entire treatment period simply as a

product of associated probabilities of realized treatment assignment at every time point. These probabilities could then be used to adjust the naive head-to-head comparison via inverse probability weighting when comparing the mean outcomes for patients who received different treatment sequences. In reality, we do not know the underlying mechanism of treatment assignment, given patient's current or past severity, but having a large sample of data with observed confounders (here, previous treatment outcomes) that may have driven physicians' treatment decisions, we can estimate probabilities of observed treatment regimes from the data using appropriate modeling techniques, e.g., logistic regression, and use these estimated probabilities to draw final inference. Therefore, MSM proceeds in two stages. At the first stage, probability of observed treatment regimes is modeled based on various observed intermediate outcomes and patient characteristics as predictor variables. Note that it is essential to have non-trivial probabilities of assignment to either treatment groups A or B, whether patients are experiencing higher or lower levels of depression severity (so that the conditional probabilities of treatment assignment, given patients current severity status are well-defined and bounded away from 0 or 1), otherwise the weights would be unstable or impossible to estimate. Such an adverse situation may occur, for example, if physicians are deterministically assigning patients whose severity score at some time point exceeds a pre-defined cutoff to treatment B, leaving no chance of observing patients with the same severity scores but receiving A. This is sometimes called the experimental treatment assumption (ETA) [27]. Also, the success of the MSM depends on the validity of the assumption of no unmeasured confounders, essentially meaning that all relevant variables that may be driving physicians' decisions are collected and properly accounted for in estimating the probability of treatment assignment. For a more technical treatment of this and other MSM assumptions, see Robins et al. 2000 [25]. At the second stage, the outcomes associated with different treatment regimes are compared using standard statistical analysis (such as ANCOVA, logistic regression, estimating equations for repeated measures or proportional hazards Cox model for time to event outcome) applied to the data where each patient's record is weighted inversely to the estimated probability of treatment regime observed for that patient. Time-varying weights can be incorporated in analyses where the model is repeated measures or time to event.

Conclusion

While traditional double-blind randomized placebo-controlled trials, in which a new approval-seeking drug is compared to placebo, are characterized by high internal validity, this type of study has substantial limitations regarding the extent to which the study results can be transferred into routine clinical care in the treatment of affective disorders. Therefore, new methodological approaches which may overcome these problems are clearly needed, with marginal structural models and Q-learning representing two of the most promising approaches in this field.

¹ It is not necessary to estimate $V(\pi)$ as it is sufficient to determine the sign of $V(\pi) - V(\pi^*)$ for each π in Π . However, we will not pursue such technicalities further (see Zhao et al. [18]; Zhang et al. [19, 21] for a discussion).

Compliance with Ethics Guidelines

Conflict of Interest

Emanuel Severus reports personal fees and other from Lundbeck, personal fees from Servier, outside the submitted work.

Eric Laber and Ilya A. Lipkovich declare that they have no conflict of interest.

Human and Animal Rights and Informed Consent

This article does not contain any studies with human or animal subjects performed by any of the authors.

References and Recommended Reading

Papers of particular interest, published recently, have been highlighted as:

- Of importance

1. Wisniewski SR, Rush AJ, Nierenberg AA, Gaynes BN, Warden D, Luther JF, et al. Can phase III trial results of antidepressant medications be generalized to clinical practice? A STAR*D report. *Am J Psychiatry*. 2009;166:599–607.
 2. Brody B, Leon AC, Kocsis JH. Antidepressant clinical trials and subject recruitment: just who are symptomatic volunteers? *Am J Psychiatry*. 2011;168:1245–7.
 3. Rutherford BR, Sneed JR, Roose SP. Does study design influence outcome? The effects of placebo control and treatment duration in antidepressant trials. *Psychother Psychosom*. 2009;78:172–81.
 4. Sinyor M, Levitt AJ, Cheung AH, Schaffer A, Kiss A, Dowlati Y, et al. Does inclusion of a placebo arm influence response to active antidepressant treatment in randomized controlled trials? Results from pooled and meta-analyses. *J Clin Psychiatry*. 2010;71:270–9.
 5. Papakostas GI, Fava M. Does the probability of receiving placebo influence clinical trial outcome? A meta-regression of double-blind, randomized clinical trials in MDD. *Eur Neuropsychopharmacol*. 2009;19:34–40.
 6. Severus E, Seemuller F, Berger M, Dittmann S, Obermeier M, Pfennig A, et al. Mirroring everyday clinical practice in clinical trial design: a new concept to improve the external validity of randomized double-blind placebo-controlled trials in the pharmacological treatment of major depression. *BMC Med*. 2012;10:67.
 7. Kelley JM, Kaptchuk TJ, Cusin C, Lipkin S, Fava M. Open-label placebo for major depressive disorder: a pilot randomized controlled trial. *Psychother Psychosom*. 2012;81:312–4.
 8. Bourin MS, Severus E, Schronen JP, Gass P, Szamosi J, Eriksson H, et al. Lithium as add-on to quetiapine XR in adult patients with acute mania: a 6-week, multi-center, double-blind, randomized, placebo-controlled study. *Int J Bipolar Disord*. 2014;2:14.
 9. • Wu F, Laber EB, Lipkovich IA, Severus E. Who will benefit from antidepressants in the acute treatment of bipolar depression? A reanalysis of the STEP-BD study by Sachs et al. 2007, using Q-learning. *Int J Bipolar Disord*. 2015;3:7.
- Clinically relevant application of Q-learning, in the context of the acute treatment of bipolar depression.
10. Nahum-Shani I, Qian M, Almirall D, Pelham WE, Gnagy B, Fabiano GA, et al. Experimental design and primary data analysis methods for comparing adaptive interventions. *Psychol Methods*. 2012;17:457–77.
 11. Nahum-Shani I, Qian M, Almirall D, Pelham WE, Gnagy B, Fabiano GA, et al. Q-learning: a data analysis method for constructing adaptive interventions. *Psychol Methods*. 2012;17:478–94.
 12. Schulte PJ, Tsiatis AA, Laber EB, Davidian M. Q- and a-learning methods for estimating optimal dynamic treatment regimes. *Stat Sci*. 2014;29:640–61.

13. Shortreed SM, Laber E, Scott ST, Pineau J, Murphy SA. A multiple imputation strategy for sequential multiple assignment randomized trials. *Stat Med*. 2014;33:4202–14.
 14. Goldberg Y, Kosorok MR. Q-learning with censored data. *Ann Stat*. 2012;40:529–60.
 15. Zhao YQ, Zeng D, Laber EB, Song R, Yuan M, Kosorok MR. Doubly robust learning for estimating individualized treatment with censored data. *Biometrika*. 2015;102:151–68.
 16. Laber EB, Lizotte DJ, Ferguson B. Set-valued dynamic treatment regimes for competing outcomes. *Biometrics*. 2014;70:53–61.
 17. Laber EB, Zhao YQ, editors. Tree-based methods for individualized treatment regimes. In: *Biometrika*. 2015 (in press).
 18. Zhao YQ, Zeng DL, Rush AJ, Kosorok MR. Estimating individualized treatment rules using outcome weighted learning. *J Am Stat Assoc*. 2012;107:1106–18.
 19. Zhang BQ, Tsiatis AA, Laber EB, Davidian M. A robust method for estimating optimal treatment regimes. *Biometrics*. 2012;68:1010–8.
 20. Zhang Y, Laber EB, Tsiatis AA, Davidian M, editors. Using decision lists to construct interpretable and parsimonious treatment regimes. In: *Biometrics*. 2015. in press
 21. Zhang B, Tsiatis AA, Davidian M, Zhang M, Laber E, editors. Estimating optimal treatment regimes from a classification perspective. In *Stat*. 2012;1:103–114.
 22. Hernán MA, Brumback B, Robins JM. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology*. 2000;11:561–70.
 23. Hernán MA, Brumback B, Robins JM. Marginal structural models to estimate the joint causal effect of nonrandomized treatments. *J Am Stat Assoc*. 2001;96:440–8.
 24. Hernán MA, Brumback BA, Robins JM. Estimating the causal effect of zidovudine on CD4 count with a marginal structural model for repeated measures. *Stat Med*. 2002;21:1689–709.
 25. Robins JM, Hernán MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology*. 2000;11:550–60.
 26. • Severus E, Lipkovich I, Seemuller F, Obermeier M, Grunze H, Bernhard B, et al. The potential role of marginal structural models (MSMs) in testing the effectiveness of antidepressants in the treatment of patients with major depression in everyday clinical practice. *World J Biol Psychiatry*. 2013;14:386–95.
- A simulation experiment used to illustrate an application of marginal structural models (MSMs) in the context of the treatment of a major depressive episode.
27. Mortimer KM, Neugebauer R, van der Laan M, Tager IB. An application of model-fitting procedures for marginal structural models. *Am J Epidemiol*. 2005;162:382–8.