Published in final edited form as:

Clin Trials. 2014 August ; 11(4): 393–399. doi:10.1177/1740774514527651.

Introduction to Dynamic Treatment Strategies and Sequential Multiple Assignment Randomization

Philip W. Lavori and
Stanford University

Ree Dawson
Frontier Science

Abstract

Background—In June 2013 a one-day workshop on Dynamic Treatment Strategies (DTS) and Sequential Multiple Assignment Randomized Trials (SMART) was held at the University of Pennsylvania in Philadelphia, PA. These two linked topics have generated a great deal of interest as researchers have recognized the importance of comparing entire strategies for managing chronic disease. A number of articles emerged from that workshop.

Purpose—The purpose of this survey of the DTS/SMART methodology (which is taken from the introductory talk in the workshop) is to provide the reader of the collected articles presented in this volume with sufficient background to appreciate the more detailed discussions in the papers.

Methods—The way that the DTS arises naturally in clinical practice is described, along with its connection to the well-known difficulties of interpreting the analysis by intention-to-treat. The SMART methodology for comparing DTS is described, and the basics of estimation and inference presented.

Results—The DTS/SMART methodology can be a flexible and practical way to optimize ongoing clinical decision making, providing evidence (based on randomization) for comparative effectiveness.

Limitations—The DTS/SMART methodology is not a solution for unstandardized study protocols.

Conclusions—The DTS/SMART methodology has growing relevance to comparative effectiveness research and the needs of the learning healthcare system.

Keywords

Clinical trials; Dynamic treatment strategies; Sequential Multiple Assignment Randomization

1 Introduction

If after initial treatment for a disease, a patient neither recovers fully nor succumbs irreversibly, then the ongoing management of the disease takes on a dynamic or adaptive character. For some time, the clinicians treating the patient will rely on changes in patient status to guide a series of treatment decisions, according to some more or less well-specified algorithm. This continues until the patient reaches the endpoint that defines the goals of treatment. A Dynamic Treatment Strategy (DTS) is a particular set of decision rules for ongoing management, with a definite format and content, which we can represent by a time-indexed mapping from the history of patient states and previous treatments into the set of possible current treatment choices.

For example, the primary care physician treating a patient for asymptomatic hypertension uses the patient's blood pressure as an indicator of the need to change the therapeutic approach to avoid stroke or heart attack. The oncologist treating a patient with lymphoma aims to prolong life, but uses intermediate events, such as tumor response to induction treatment or subsequent progression of disease to guide the use of second-line therapy or maintenance treatments. A clinic has several initial approaches to weight-loss, including diet and exercise, but may need to change the approach to medication or surgery if the patient is not succeeding.

Clinicians are now faced with the task of developing effective DTS, testing them for efficacy, and comparing their effectiveness. As more treatment options are developed for a clinical condition, the opportunities for sequencing them increase, with a consequent need to understand how they can be used to best effect. The generalization of the conventional trial suitable for the comparison of such DTS is called the Sequential Multiple Assignment Randomized Trial (SMART). For example, in the B-cell lymphoma trial described in [1], patients were randomized to induction treatment with either CHOP or CHOP+rituximab (R-CHOP). For patients whose tumors do not respond to the initial treatment (either CHOP or R-CHOP) a standard (supportive) care was offered but in those whose tumor did respond, the next clinical decision involved a choice between observation (O) or maintenance with rituximab (MR), regardless of what induction treatment produced the tumor response. Thus, the study design called for a second-stage randomization of the maintenance treatment in patients whose tumor responded to the induction treatment assigned in the first randomization.

This SMART design thus implicitly defines four DTS, each one a complete strategy for the patients entering the study, determined by the induction treatment and the maintenance treatment if induction produced a response. Note that each of these DTS has a branching structure with one branch point that is defined by the intermediate outcome (tumor response) to the initial treatment (induction). The DTS differ from each other on the choice of what to do for induction and then what to do for maintenance in responders. The trial adds a second kind of branching (the randomization among treatments). It is useful to distinguish between the two kinds of branching, since the whole SMART can be represented as a big tree, with layers of clinical state branching (defining the DTS) and treatment branching (the randomizations).

The June 2013 conference in Philadelphia brought together a diverse group of clinicians, statisticians, regulators, drug developers, and others, to discuss methodological issues relating to that task. In this introduction to the more specific and technical articles arising from the Philadelphia conference we set up the basic definitions and general structure of the task, present a few historical notes on its origins, and compare the SMART methodology to the various less general and more conventional approaches that have been proposed or used. We also describe the two complementary views of the role of the SMART methodology: for exploratory development of a new, optimized DTS, which would then be tested in a conventional trial comparing it to a control treatment, and for confirmatory comparison of a small number of *a priori* specified DTS. We briefly allude to the technical and practical issues that come up, which are covered in more detail in the subsequent papers.

The intended audience for this paper includes a broad mixture of clinical investigators, reviewers and designers of protocols, and regulators. We include material that will inform the more clinically inclined portion of the audience, because the uptake of these methods is limited by unfamiliarity among clinical investigators and grant reviewers. We show that the DTS formalism arises naturally and inevitably from clinical considerations. We illustrate the flexibility and practical utility of the SMART methodology, by showing how it responds to difficulties with the intention-to-treat (ITT) principle of clinical trial design and analysis. Our intent is to demystify the subject, and provide a gentle introduction that may be of use to grant proposal and manuscript reviewers as well as investigators, given the formidable barriers our conservative clinical trials field erects in the path of design innovation. To be blunt, it seems necessary to show that SMART methodology is not some radical departure from the conventional randomized trial. It does not involve instrumental variables, complex modeling of covariates to adjust for confounding by indication, nor does it appeal to “Big Data” magic.

2 Origins

Early references to DTS, and the equivalent (and widely used) “Adaptive Treatment Strategy”, “Dynamic Treatment Regime”, or ‘treatment policy’, include [2, 3, 4, 5, 6, 7, 8]. We use DTS here to avoid the collision of meaning between “adaptive trial” (the methodological concept) and “adaptive strategy” (the clinical concept), which has been a source of confusion. In addition, the DTS terminology reminds us of the original usage in “dynamic control”, which is the much older engineering and operations research context in which the key ideas were first defined. The term “regime” seems to conjure up a sinister political entity for some; the connection to the medical term “regimen” is apparently somewhat obscure.

In the mid 1980’s, a few cancer research groups designed and conducted rigorous studies to evaluate 2-stage DTS with survival outcomes [9]. Another independent line was pursued in the late 1990s in psychiatry, and the general K-stage DTS was described in [2, 4, 9]. In the 2000s DTS were intensively studied for use in substance abuse research [10]. These DTS were first conceived as a representation of what clinicians generally do to manage the treatment of patients. Meanwhile, epidemiologists were developing methods motivated by the following generic epidemiologic problem: Given a sample of subjects with time-varying

exposure, how can one estimate the consequences of continuous exposure? Robins developed a longitudinal generalization [11] of Rubin's potential outcomes framework [12] for inference about treatment/exposure effects from observational data. The methods were soon generalized to the question of estimating the results of a DTS for disease management from observational data on naturalistic treatment. Robins provided a full development of methods for estimating and comparing outcomes across all DTS represented in the data, beginning with G-estimation and culminating in the methods based on optimal semi-parametric theory [13], and methods designed to cope with the combinatorial explosion of observed DTS in naturalistic data. A comprehensive review is in [14].

At the same time, experimental evaluation of DTS was developed, motivated by the perennial problem of unobserved confounding in observational studies for effects of treatments. By extension of the 1-stage treatment context, the key assumption is that (at each stage) the treatments are assigned in a way that is conditionally independent of the potential future responses to any future treatment sequence, given the history of states and treatments up to the current stage ("sequential ignorability").

Confounding by indication (failure of ignorability) is a much less tractable problem in the treatment context than in the exposure context. People seldom choose their environmental exposures on the basis of their consequences to health. But in the treatment context a dedicated, highly trained cadre of individuals, called clinicians, choose treatments *precisely* on that basis, and thus virtually guarantee confounding. The combinatorial explosion of variation in naturalistically-selected treatments also interferes with rigorous evaluation; too many possible strategies represented in the data, with too little rationale for the variation.

Experimentalists have used randomized clinical trials to evaluate and compare DTS, most commonly by randomizing patients once at baseline to a small number (usually only two) competing DTS [15], and then directly comparing outcomes as usual. More recently, investigators have used sequential randomization instead of baseline randomization. Sequential randomization makes the assumption of sequential ignorability true by design. Thus all the methods developed for the observational case are applicable. In the next section we show how consideration of the woes of interpretation of the ITT principle motivate the Sequential Multiple Assignment Randomized Trial (SMART), with examples, discuss estimation and inference from such a trials and compare the efficiency of baseline randomized and SMART designs.

3 From ITT to SMART

Many clinicians, including sophisticated investigators, find ITT difficult to swallow because it ignores the changes in treatment (from the originally randomized treatment) that the patient undergoes in the intervening period between randomization and outcome. One can sometimes argue that ITT gets at something useful even in the face of non-adherence, because it describes the effects of policies (try to give A versus try to give B) insofar as the non-adherence in a trial reflects what will happen in practice. But when the changes in treatment after randomization are clinically motivated and strongly indicated, the clinician becomes (rightly) suspicious of the policy rationale. The ITT analysis appears to ignore the

fact that the clinicians treating the patient in the trial have to follow a DTS (even though they may not think of it that way.)

For example, consider the varied clinical interpretations offered in response to the (ITT) analyses of one of the pivotal trials comparing paclitaxel (P) vs paclitaxel + bevacizumab (P +B) in advanced breast cancer [16]. That trial showed an advantage of the combination in lengthening the time to progression (TTP), which was the primary outcome, but did not show an effect on overall survival time (OS). Some reviewers (on the FDA panel and elsewhere) expressed a preference for the OS analysis, discounting the clinical importance of TTP, while others preferred TTP because in their view, the changes in treatment that happen when a patient's cancer progresses preclude interpretation of the OS difference (or lack thereof). One reviewer used the phrase "muddies the waters" in describing the effects of these subsequent rescue treatments. The water is considerably clearer after recasting the problem in terms of DTS.

Ignoring all other sources of change, we will consider only what happens as a result of progression, and imagine that the treating clinician has in mind what should be done if the patient experiences disease progression. Each choice of treatment for the progression turns the original assignment (say, to P+B) into a DTS. For a given physician, the decision among alternative treatments for progression may depend on other aspects of the patient's state at progression, or the clinician may think there is only one choice for all. Two clinicians might differ in the way they deal with progression. In general, there are several hidden strategies behind the scenes in the trial, each of which contributes to the "muddying of the water" for the ITT analysis of OS. In this sense, the ITT analysis averages over the hidden mixture of DTS.

How can the DTS idea help interpret the ITT analysis? We can describe two extreme cases:

1. Suppose all the DTS that are observed to extend P and P+B are optimal, in the obvious sense that no patient receives a treatment that would result in a shorter survival after progression than if he or she had received a different treatment (from among the available options). Note that the optimal treatment for progression after P+B may not be the same as for progression after P. Then the ITT comparison gives a good answer, in fact it is ideal for the purposes of determining how the overall optimal DTS should start. There is not much mud in the water! One can substitute "nearly all DTS" and "close to optimal" and even begin to model how much it would matter. Note that the optimality condition also holds if progression leads quickly and inexorably to death regardless of treatment after progression, so that all the extensions of P and P+B are equally ineffective.
2. Suppose there are popular DTSs that reveal disagreement about what to do with progression after P or P+B, and these varying decisions have adverse consequences for OS (that is, are not all close to optimal). Then the results of a study will misinform us about the prognostic significance (for survival) of progression on P and progression on P+B, because some clinicians make suboptimal choices of treatment after progression on one

or the other (or both). There is lots of mud in the water here, but the mud also affects the value of the TTP analysis as a reliable proxy for survival. The substantial uncontrolled and unobservable variation in the optimality of the progression treatments makes the TTP comparison unhelpful unless investigators can assert that the TTP difference is reliable for the survival difference on optimal post-progression treatment, and it is hard to see how we would know that in case (2). We live in a DTS clinical world, but we do conventional trials whose ITT interpretations depend on tacit assumptions about the optimality of the DTS extensions.

One route to SMART begins with the kind of design discussion described above, in the process of planning a conventional randomized trial. The planners identify a crucial intermediate outcome (such as progression) between the randomization and the primary endpoint (such as survival). They attempt to address the methodological problem caused by the clinical need to intervene in reaction to the intermediate outcome, which intervention they worry will “muddy the waters” of the primary ITT analysis. An attempt may be made to standardize those interventions by protocol, at least to confine variation to a set of choices that can be defended as near optimal. If they succeed, the conventional trial is strengthened to the degree that the remaining DTS variation does not stray far from case (1) above.

But it may be that there is genuine uncertainty about how to intervene, with some arguing for intervention X after progression on one of the initial treatments and some arguing for Y, and there may be different options X' and Y' for intervention after progression on the other initial treatment. So the attempt at standardization fails, leaving the conventional trial without a clear way forward, other than to give up on OS and change the primary endpoint to progression. The SMART methodology provides a solution: randomize the progression interventions, too.

4 SMART for trials, DTS for patients

It is important to bear in mind that the DTS is a *treatment* strategy, while the SMART design is an *experiment* for comparing among DTS. A single DTS as a clinical object (something that doctors do) is a tree with only state-branching. There is no branching in the treatments for patients with a particular history of states and previous treatments. This is another way of stating that a single DTS calls for an unambiguous, single valued treatment depending on the patient's history of previous treatments and the resulting clinical states. A SMART is a tree with both kinds of branching because it compares among several DTS, by randomizing patients to the different treatment branches that separate the different DTS. A conventional single stage trial may also have both kinds of branches, but the treatment branching after the first (randomized) branching is not randomized, and the ITT analysis mixes over the distribution of the implied DTS as discussed above. Randomizing the treatments at occasions where the patient's clinicians might ordinarily be the ones driving the differences in treatments, converts a mixed DTS ITT trial into a SMART. We hasten to add that randomizing the primary treatment branches may not remove all variation in the treatment regimes. Thus, the SMART design is not a solution for unstandardized study protocols, but

rather a tool for addressing important open questions that hinder the construction (and implementation) of an optimal DTS.

The Sequential Multiple Assignment Randomized Trial guarantees “sequential ignorability”, which generalizes Rubin’s concept of ignorability to the DTS context. We suppress some technical matters here, see [17] for a more thorough discussion. With that guarantee, the observational data methods can be borrowed to provide several possible approaches to estimation and testing of the results of different DTS represented in the SMART design. The three main approaches are G-estimation [18], marginal means (MM) [19, 8, 5], and the so-called optimal semi-parametric estimator (OSP). Moodie et al. [20] evaluate alternative approaches to inference from SMAR trial data, and Nahum-Shani et al. [21] present a comprehensive discussion of design and analysis, with details of computation, oriented toward the behavioral scientist.

Under mild assumptions, the the G and MM estimators are just differently weighted averages of the fully observed stratum means [22]. and the G and OSP estimators coincide if the observed and expected randomization proportions coincide. For different purposes it is easier to work with one or the other, providing another reason to use forced balancing of randomization to push the observed assignment rates toward the expected ones by design.

Several approaches to variance estimation for the G-estimate all yield the same variance estimate [22]. In addition, when the observed and expected assignments to treatment coincide, the optimal semi-parametric variance estimate also agrees (see [23]).

The sequential randomization implies that a patient may contribute data to the estimates of more than one DTS; indeed, every patient contributes data toward a DTS until that patient departs from the DTS (by a randomization). The estimates of the expected outcomes for those two DTS are therefore in general correlated. See [22, 17] for details. The covariance of the SP estimators of two DTS is discussed in [17].

The principal alternative to SMART is baseline randomization among the various DTS under study. The methods described above can be used for inferences in either design, although one would ordinarily analyze data from baseline randomization in the usual way, ignoring the structure of the DTS, since the treatment groups are non-overlapping. There have been various statistical attempts to describe the relative inefficiency of baseline randomization [24, 25] compared to SMART.

5 Two roles for SMART

There are two differing roles for SMART designs: (1) to compare a small number of DTS that are well-specified *a priori* and hence are embedded in the SMART, and (2) to explore the possibility of optimizing a new DTS, possibly beyond those embedded in the design. Defining the states and the treatment options that will be considered is at the heart of the design of a SMART. In the 2-stage cancer trials mentioned above, there are a small number of well-defined options, and the intent may be to conduct a confirmatory, Phase III trial of the resulting small set of DTS. But in many contexts, the potential for tailoring the treatments to the patient’s changing state may be much less well specified *a priori*. Murphy

([19]) argues that one should “...restrict the class of treatment options at each decision point only by ethical, scientific, or feasibility considerations”, which are captured in the states, which are in Murphy’s view only to “involve information that is to be used to constrain treatment options”, and not other auxiliary information. On this view, including the auxiliary information in the states implies that the strategies under consideration must treat patients with different auxiliary information (but the same state) differently, in effect limiting the strategies being considered. By using only the state information, and including a wide range of treatment options, the auxiliary information can be used in secondary analyses to see if they should be used to discriminate between treatments (see below for a brief discussion of Q-learning). Murphy also proposes that the SMART design should be used for “one trial in a series of developmental, randomized trials...prior to a confirmatory trial”. The developmental trials might use SMART methods to find a candidate optimal DTS (see below), which would then be tested in a conventional confirmatory trial against the standard treatment. SMART developmental designs can also be used to explore the optimal way to use “tailoring variables”, which are the measures of patient state used in the rules to assign treatment. For example, if the patient’s subsequent treatment is to depend on response to a initial treatment, randomizing the choice of cutpoint on the symptom scale that defines response allows comparison of DTS’s that differ on that component rather than on the treatments themselves, which may be well-accepted already. Timing of changes, such as how long a patient’s initial remission should extend before the clinician considers switching to a maintenance treatment) can also be randomized.

6 Optimizing the DTS: Q-Learning

So far we have emphasized the simplest kind of analysis of data from a SMART, where the common scaffold of the DTS (particularly the states used to drive the treatment decisions as well as the rules themselves) are fixed in advance, allowing the planned use of SMART to compare the DTS. However, investigators can also use SMART methodology as a development tool, to construct a putative optimal DTS, which then may be tested in a conventional “confirmatory” head-to-head trial against a control treatment (which would in general not be adaptive). The intuition is similar to the idea of using baseline factors in a conventional randomized trial of treatments in an exploratory statistical analysis to develop tailored or personalized treatment decisions. That familiar idea, taken to the multi-stage context of a SMART, turns into Q-learning. The article by Nahum-Shani [26] is an accessible introduction to Q-learning in this context, and the new text by Chakraborty and Moodie [27] has a comprehensive discussion of the method and its difficulties.

7 Non-adherence and DTS

Another example of a DTS in practice arises as a solution to the common clinical problem of non-adherence. A patient with hypertension may decide not to take the prescribed blood pressure medication, calling for an adaptive action by the clinician. Assuming it would not be acceptable for the doctor to stand by and wait for a stroke or heart attack, the clinical choices might depend on the reasons for this patient’s desire to stop. Unacceptable side-effects suggest a change of medications, while a less specific reason (hypertension is asymptomatic, and the patient loses motivation over time) might suggest a reminder of why

maintaining lower blood pressure is important. Adherence problems call for adaptive responses. Behavioral interventions (diet, exercise, smoking cessation, and others) tend to elaborate naturally into DTS. We have had several recent consulting experiences involving the application of SMART methodology for studying DTS with adherence interventions in, for example, smoking cessation and management of Type 1 diabetes in children. Although reports of these trials have yet to appear in the literature, we note that references to the idea of using non-adherence as a tailoring variable or target of intervention have appeared [28, 29].

8 Adaptive trial designs and SMART

In a DTS, the treatments are adaptive to the changes in the patient's state. By contrast, the "conventional" SMART has a fixed design that does not "adapt" to outcomes in patients. Notably, Thall [9, 30] used trial adaptation along with SMART to drop non-performing treatments. It seems likely that other adaptive methods could be applied to SMART design. For example, Shih and Lavori [31], point out that the combinatorial expansion of DTS that can occur in a SMART makes it attractive to consider ways to abandon entire DTS early in the trial so that sample size can be concentrated on winning strategies. The general area seems ripe for further development.

Acknowledgments

The authors acknowledge support from grant R01 MH051481 to Stanford University from the National Institute of Mental Health and PWL is also supported by P30 CA124435 to Stanford from the National Cancer Institute and by the Clinical and Translational Science Award UL1 RR025744 to the Stanford Center for Clinical and Translational Education and Research (Spectrum) from the National Center for Research Resources, National Institutes of Health to Stanford University.

References

1. Habermann TM, Weller EA, Morrison VA, Gascoyne RD, Cassileth PA, Cohn JB, Shaker RD, Woda B, Fisher RI, Peterson BA, Horning SJ. Rituximab-chop versus chop alone or with maintenance rituximab in older patients with diffuse large b-cell lymphoma. *Journal of Clinical Oncology*. 2006; 24:3121–3127. [PubMed: 16754935]
2. Lavori PW, Dawson R. Developing and comparing treatment strategies: An annotated portfolio of designs. *Psychopharmacological Bulletin*. 1998; 34(1):13–18.
3. Lavori PW, Dawson R, Rush AJ. Flexible treatment strategies in chronic disease: Clinical and research implications. *Biological Psychiatry*. 2000; 48:605–614. [PubMed: 11018231]
4. Lavori PW, Dawson R. A design for testing clinical strategies: biased individually tailored within-subject randomization. *Journal of the Royal Statistical Society, Series A*. 2000; 163:29–38.
5. Murphy SA, VanderLaan MJ, Robins JM. Marginal mean models for dynamic treatment regimes. *Journal of the American Statistical Association*. 2001; 96(456):1410–1423. [PubMed: 20019887]
6. Lunceford JK, Davidian M, Tsiatis AA. Estimation of survival distributions of treatment policies in two-stage randomization designs in clinical trials. *Biometrics*. 2002; 58:48–57. [PubMed: 11890326]
7. Lavori PW, Dawson R. Dynamic treatment regimes: practical design considerations. *Clinical Trials*. 2003; 1:9–20.
8. Murphy SA. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society Series B*. 2003; 65:531–66.
9. Thall P, Millikan R, Sung HG. Evaluating multiple treatment courses in clinical trials. *Statistics in Medicine*. 2000; 19:1011–1028. [PubMed: 10790677]

10. Murphy SA, Lynch KG, Oslin D, McKay JR, TenHave T. Developing adaptive treatment strategies in substance abuse research. *Drug and alcohol dependence*. 2007; 88:S24–30. [PubMed: 17056207]
11. Robins JM. A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect. *Mathematical Modeling*. 1986; 7:1393–1512.
12. Holland PW. Statistics and causal inference. *JASA*. 1986; 81
13. Robins JM, Rotnitzsky A, Zhao LP. Estimation of regression coefficients when some regressors are not always observed. *JASA*. 1994; 89:846–866.
14. Robins, JM.; VanDerLaan, J. Unified methods for censored longitudinal data and causality. Springer; 2010.
15. Boden WE, ORourke RA, Crawford MH, Blaustein AS, Deedwania PC, Zoble RG, Wexler LF, Kleiger RE, Pepine CJ, Ferry DR, Chow BK, Lavori PW. Outcomes in patients with acute non-q-wave myocardial infarction randomly assigned to an invasive as compared with a conservative strategy. *New England Journal of Medicine*. 1998; 338(25):1785–1792. [PubMed: 9632444]
16. Miller K, Wang M, Gralow J, Dickler M, Cobleigh M, Perez EA, Shenkier T, Cella D, Davidson NE. Paclitaxel plus bevacizumab versus paclitaxel alone for metastatic breast cancer. *New England Journal of Medicine*. 2007; 357(26):2666–2676. [PubMed: 18160686]
17. Dawson R, Lavori PW. Sequential causal inference: Application to randomized trials of adaptive treatment strategies. *Statistics in Medicine*. 2008; 34(27):1626–45.
18. Robins JM. The control of confounding by intermediate variables. *Statistics in Medicine*. 1989; 8:679–701. [PubMed: 2749074]
19. Murphy SA. An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine*. 2005; 24:1455–1481. [PubMed: 15586395]
20. Moodie EEM, Richardson TS, Stephens DA. Demystifying optimal dynamic treatment regimes. *Biometrics*. 2007; 63:447–55. [PubMed: 17688497]
21. Nahum-Shani I, Qian M, Almirall D, Pelham WE, Gnagy B, Fabiono GA, Waxmonsky JG, Jihnhee Yu, Murphy SA. Experimental design and primary data analysis methods for comparing adaptive interventions. *Psychological Methods*. 2012; 17:457–477. [PubMed: 23025433]
22. Lavori PW, Dawson R. Improving the efficiency of estimation in randomized trials of adaptive treatment strategies. *Clinical Trials*. 2007; 4:297–308. [PubMed: 17848491]
23. Dawson R, Lavori PW. Efficient design and inference for multistage randomized trials of individualized treatment policies. *Biostatistics*. 2011; 1:1–10.
24. Dawson R, Lavori P. Comparison of designs for adaptive treatment strategies. *J Stat Planning and Inference*. 2003; 117:365–385.
25. Wahed AS, Tsiatis AA. Optimal estimator for the survival distribution and related quantities for treatment policies in two-stage randomization designs in clinical trials. *Biometrics*. 2004; 60:124–133. [PubMed: 15032782]
26. Nahum-Shani I, Qian M, Almirall D, Pelham WE, Gnagy B, Fabiono GA, Waxmonsky JG, Jihnhee Yu, Murphy SA. Q-learning: A data analysis method for constructing adaptive interventions. *Psychological Methods*. 2012; 17:478–494. [PubMed: 23025434]
27. Chakraborty, B.; Moodie, EM. Statistical methods for dynamic treatment regimes. Springer; 2013.
28. Murphy SA, Oslin DW, Rush AJ, Zhu J. Methodological challenges in constructing effective treatment sequences for chronic psychiatric disorders. *Neuropsychopharmacology*. 2007; 32:257–62. [PubMed: 17091129]
29. Buyze J, Van Rompaye B, Goetghebeur E. Designing a sequentially randomized study with adherence enhancing interventions for diabetes patients. *Statistics in Medicine*. 2010; 29:1114–26. [PubMed: 20101597]
30. Thall PF, Sung HG, Estey EH. Selecting therapeutic strategies based on efficacy and death in multicourse clinical trials. *JASA*. 2002; 97:29–39.
31. Shih MC, Lavori PW. Sequential methods for comparative effectiveness experiments: point of care clinical trials. *Statistica Sinica*. 2013; 23:1775–1791.