Published in final edited form as:

Stat Med. 2008 May 10; 27(10): 1626–1645. doi:10.1002/sim.3039.

Sequential causal inference: Application to randomized trials of adaptive treatment strategies

Ree Dawson^{†,1,*} and Philip W. Lavori²

¹ Frontier Science & Technology Research Foundation, 900 Commonwealth Ave., Boston, MA 02215, U.S.A

² Department of Health Research and Policy, M/C 5405, Stanford, CA 94305, U.S.A

SUMMARY

Clinical trials that randomize subjects to decision algorithms, which adapt treatments over time according to individual response, have gained considerable interest as investigators seek designs that directly inform clinical decision making. We consider designs in which subjects are randomized sequentially at decision points, among adaptive treatment options under evaluation. We present a sequential method to estimate the comparative effects of the randomized adaptive treatments, which are formalized as adaptive treatment strategies. Our causal estimators are derived using Bayesian predictive inference. We use analytical and empirical calculations to compare the predictive estimators to (i) the ‘standard’ approach that allocates the sequentially obtained data to separate strategy-specific groups as would arise from randomizing subjects at baseline; (ii) the semi-parametric approach of marginal mean models that, under appropriate experimental conditions, provides the same sequential estimator of causal differences as the proposed approach. Simulation studies demonstrate that sequential causal inference offers substantial efficiency gains over the standard approach to comparing treatments, because the predictive estimators can take advantage of the monotone structure of shared data among adaptive strategies. We further demonstrate that the semi-parametric asymptotic variances, which are marginal ‘one-step’ estimators, may exhibit significant bias, in contrast to the predictive variances. We show that the conditions under which the sequential method is attractive relative to the other two approaches are those most likely to occur in real studies.

Keywords

adaptive treatment strategy; sequential causal inference; multi-stage decision problems; sequential randomization; individualized treatments

1. INTRODUCTION

Statisticians have formalized clinical decision making as adaptive treatment strategies or dynamic regimes in order to develop experimental trial methodology that is of direct, practical value to clinicians [1,2]. Conceptually, adaptive treatment strategies are dynamic algorithms for matching clinical decisions to the evolving course of the individual patient’s response to treatment over time. By their very nature, adaptive strategies are carried out as sequences of decision rules that give rise to sequences of treatments, in contrast to the fixed treatments that

*Correspondence to: Ree Dawson, Frontier Science & Technology Research Foundation, 900 Commonwealth Ave., Boston, MA 02215, U.S.A.

[†]E-mail: dawson@fstrf.dfci.harvard.edu

are classically studied in randomized controlled trials (RCTs). Clinical trials for evaluating adaptive strategies are already in progress in the areas of mental health disorders, rapidly fatal diseases, and prevention programs [3–7]. Designs have also been developed to evaluate the clinical decision making that underlies the sequential monotherapy or polytherapy regimes characteristic of mental health treatments [8–11].

A typifying example of an adaptive strategy occurs in the treatment of a chronic disorder such as depression. Consider patients with unipolar depression who do not respond well to medication, and assume that subjects are started on a standard antidepressant. After 8 weeks on initial treatment, subjects are rated in terms of their treatment response: none, partial, or full. A particular adaptive strategy then specifies a second-line option for patients who fail to remit under the standard front line, switching non-responders to an alternative antidepressant, and augmenting the treatment of partial responders. After another 8 weeks, subjects are re-assessed; the strategy specifies treatments for switching or augmenting medication of the non-remitters according to their degree of response. This strategy mimics the sequential decision making of clinicians who are faced with a ‘treatment-resistant’ patient, which often leads to adaptively determined polytherapy treatment.

The issue of how to allocate subjects to adaptive strategies in a RCT has been examined previously, from the perspective of comparing baseline with sequential randomization (SR) [1,12], and of determining allocation probabilities that are applied sequentially [2]. Here, SR means that subjects are randomized at decision points, among available treatment options specified by the strategies being evaluated. For example, consider the simplest RCT where there are two strategies under evaluation: ‘Start on the standard treatment STD; after 8 weeks, switch non-responders to SW1_i, augment treatment of partial responders with AUG1; after another 8 weeks, switch non-responders to SW2_i, augment treatment of partial responders with AUG2_i’, $i = 1, 2$. A subject with depression who fails to respond to the standard treatment will be adaptively switched to a second-line option, using randomization to assign either SW1₁ or SW1₂; in the event of continued non-response, randomization at the first decision point determines the third-line option (SW2₁ or SW2₂). Subjects who partially respond to STD, but fail to fully remit after AUG1 is added to the standard treatment, will be randomized at the second decision point to one of the two augmentation options, AUG2₁ or AUG2₂. In contrast, subjects who obtain a remission will remain on the standard treatment; it is not necessary to randomize them during the trial. This paper takes up the basic issue of inference for the comparative effects of sequentially randomized adaptive treatments.

The resistant depression example illustrates a key property of data obtained *via* SR that motivates our proposed approach to causal inference, namely the potential for there to be partially or completely overlapping data among strategies under evaluation. Shared data may occur because subjects are randomized at decision points, rather than conventionally at baseline into separate strategy-specific groups. For the simple RCT described above, strategies share data on individual responses to STD under SR (but not under baseline randomization); initial responses are used by both strategies to determine second-line treatments prior to any (sequential) randomization. Complete treatment and response data of subjects who fully remit under standard treatment are also shared by both strategies since these subjects are never randomized. The degree of overlap in SR data is determined in part by the algorithms specified by the strategies (the common treatment option AUG1 in the example RCT allows sharing of partial responder data for first- and second-line treatments), but also by the study population given the adaptive nature of the algorithms (the proportion of remitters to standard treatment determines the extent of complete overlap). In any case, the structure of shared data is always nested or monotone due to the specification of strategies as sequences of decision rules, which in turn are functions of individual patient history. Consequently, once strategies specify distinct

treatments for subjects with a certain treatment and response history, their SR data are distinct after that decision point for those subjects.

We present a *sequential* approach to inference that capitalizes on the monotone structure of shared SR data, the goal being to efficiently estimate pairwise differences in the response means of strategies under evaluation. The basic idea is to reduce the overall variance for a two-way comparison by increasing the precision for strategy-specific estimates, while taking into account between-strategy covariance due to shared data values, with focus on the intention-to-treat (IT) causal estimand. Previously, we showed that SR data can be allocated to separate groups, one per strategy, as would arise from randomizing subjects at baseline, thus allowing standard methods to be used for valid estimation of causal effects [12]. We investigate analytically and through simulation the potential efficiency gains accrued from the use of shared data that the sequential method enjoys. A key determinant of relative efficiency is how quickly strategies diverge, in the sense of adaptively specifying distinct treatments, as this governs overlap in SR data, which in turn drives the magnitude of within-strategy standard errors and between-strategy covariance.

We further consider the semi-parametric approach to inference from SR data that has been applied to the design of randomized trials of adaptive strategies [2]; under appropriate experimental conditions, our proposed sequential IT estimator coincides with that derived from marginal mean (MM) models. We show that the semi-parametric asymptotic variances for strategy-specific estimates may exhibit significant bias, in contrast to the variance estimators we develop. Additionally, the semi-parametric approach to causal inference in [2] does not take into account partially overlapping SR data among strategies, which may introduce further bias into pairwise comparisons. Such biases deserve consideration because of their effect on sample size determination.

In the next section, we set up the notation and assumptions. Then, before comparing sequential causal inference (SCI) with standard methods, referred to here as baseline inference (BI), and with the MM semi-parametric approach, we formalize the sequential inferential framework.

2. NOTATION FOR ADAPTIVE TREATMENT STRATEGIES

We consider estimation of the comparative effects of two adaptive treatment strategies ($z = c, d$) under the assumption of sequential ignorability, with emphasis on causal inference for pairwise effects arising from SR data. Suppose that assignment of a subject to strategy z gives rise to the sequence of treatments and responses:

$$S_1, A_{z,1}, \dots, S_{z,K}, A_{z,K}, Y_z$$

where S_1 is baseline or initial state (measured prior to the first treatment decision), $A_{z,k}$ are treatments, $S_{z,k}$ are post-treatment states (measured after treatment decisions), and Y_z is the final outcome, $z = c, d$.

Treatment is construed generally to include anything under the control of the investigator or clinician (e.g. choice of diagnostic assessment, choice of medication). Sequential ignorability implies that treatment $A_{z,k}$ is conditionally independent of unobserved responses given prior observed treatments and responses. We refer to $S_{z,k}$, $k > 1$, and Y_z as intermediate and final ‘potential outcomes’, respectively, noting that each subject could be potentially randomized to receive either strategy, with the corresponding potential outcomes being those observed when assigned to a particular strategy. In contrast, the sequence of treatments represents the intervention determined by the adaptive rules underlying the strategy. More precisely, the strategy $z = c$ (or $z = d$) represents a sequence of decision rules, c_k , $k = 1, \dots, K$, where $A_{c,k}$ is

determined according to c_k based on $\underline{S}_{c,k} = \{S_1, S_{c,2}, \dots, S_{c,k}\}$ and $\underline{A}_{c,k-1} = \{A_{c,1}, \dots, A_{c,k-1}\}$. The notation $A_{c,k}$ signifies determination by the decision rule c_k . State values through k , which include the values of responses to $\underline{A}_{c,k-1}$, are denoted $\underline{s}_k = \{s_1, \dots, s_k\}$, where $S_1 = s_1$ and $S_{c,j} = s_j, j > 1$.

3. SEQUENTIAL CAUSAL INFERENCE FOR PAIRWISE COMPARISONS OF ADAPTIVE TREATMENT STRATEGIES

Previously, we applied rules of probability that relate conditional means and variances to unconditional ones to derive sequential estimators for the mean and variance of an adaptive treatment strategy [13]. In this paper, we justify those calculations using a predictive (Bayesian) approach, which generalizes the results of Little and Rubin [14] for cross-sectional data (single time period) to the longitudinal (sequential) case. We also extend the sequential approach to causal inference for the average response difference of two strategies, as might be of interest for two-way comparisons based on SR data. The predictive approach provides a way to be precise about the role of correlated outcomes on causal inference.

3.1. Definition of the IT causal estimand and derivation of the predictive estimators

For a large population (of size N), define the IT causal estimand for strategies c and d as the average response difference:

$$\bar{Y}_c - \bar{Y}_d = N^{-1} \left\{ \sum_{i=1}^N Y_{c,i} - \sum_{i=1}^N Y_{d,i} \right\} \quad (1)$$

where the $Y_{z,i}$ are the individual potential final responses to strategy z . Let the population mean and proportion for the subgroup having the sequence of states $\{s_1, \dots, s_K\}$ under strategy z be, respectively:

$$\bar{Y}_z(\underline{s}_K) = N_{z,K}(\underline{s}_K)^{-1} \sum_{i=1}^N I(\underline{S}_{z,K,i} = \underline{s}_K) Y_{z,i}, \quad P_{z,K}(\underline{s}_K) = N^{-1} N_{z,K}(\underline{s}_K)$$

where $N_{z,K}(\underline{s}_K)$ is the number of subjects with $\underline{S}_{z,K} = \underline{s}_K$ and I is the indicator function of its argument. The IT causal estimand can accordingly be re-expressed as the difference in stratified population means:

$$\sum_{\underline{s}_K} P_{c,K}(\underline{s}_K) \bar{Y}_c(\underline{s}_K) - \sum_{\underline{s}_K} P_{d,K}(\underline{s}_K) \bar{Y}_d(\underline{s}_K) \quad (2)$$

where strata correspond to subgroups defined by $\underline{S}_{z,K} = \underline{s}_K$, for $z = c, d$.

The nested structure of intermediate outcomes permits writing population quantities sequentially, which we take advantage of in deriving our estimators. For $k = 1, \dots, K$, we express the subgroup proportions $P_{z,k}(\underline{s}_k) = N^{-1} N_{z,k}(\underline{s}_k)$ as the cumulative product of (conditional) response rates for $S_{z,j}, j = 1 \dots k$:

$$P_{z,k}(\underline{s}_k) = \prod_{j=1}^k F_{z,j}(\underline{s}_j) \quad \text{where } F_{z,j}(\underline{s}_j) = N_{z,j}(\underline{s}_j) / N_{z,j-1}(\underline{s}_{j-1}) \quad (3)$$

and

$$N_{z,j}(\underline{s}_j) = \sum_{i=1}^N I(\underline{s}_{z,i} = \underline{s}_j)$$

with special cases $N_{z,1}(\underline{s}_1) \equiv N_1(s_1)$, $N_{z,0}(\underline{s}_0) \equiv N$, $F_{z,1}(\underline{s}_1) \equiv F_1(s_1) = N_1(s_1)/N$. In (3) and through-out this paper, whenever two arguments $\underline{s}_j$ and $\underline{s}_{j-1}$ appear in the same expression, we always mean that the shorter sequence of states is an initial segment of the longer one.

Case $K = 1$: Suppose Y_z has a stratified normal distribution across subgroups indexed by the discrete (multinomial) values of S_1 . Let n be the sample size and $n_{z,1}(s_1)$ be the number of subjects randomized to strategy z in the subgroup with baseline value s_1 . Using the relationships for iterative expectation and variance decomposition, respectively:

$$E(A) = E[E(A|B)] \quad (4a)$$

$$V(A) = V[E(A|B)] + E[V(A|B)] \quad (4b)$$

Little and Rubin [14] derived the (posterior) predictive mean of $\bar{Y}_z$, conditional on the observed Y_z responses and S_1 values, as

$$\sum_{s_1} f_1(s_1) m_{z,1}(s_1) \quad (5)$$

and the (posterior) predictive variance of $\bar{Y}_z$ as:

$$\sum_{s_1} f_1^2(s_1) s_{z,1}^2(s_1) / n_{z,1}(s_1) + \sum_{s_1} n^{-1} f_1(s_1) \{1 - f_1(s_1)\} s_{z,1}^2(s_1) / n_{z,1}(s_1) \quad (6a)$$

$$+ \sum_{s_1, s'_1} m_{z,1}(s_1) m_{z,1}(s'_1) [n^{-1} \{I(s_1 = s'_1) f_1(s_1) - f_1(s_1) f_1(s'_1)\}] \quad (6b)$$

where $f_1(s_1)$ is the sample proportion for the subgroup defined by $S_1 = s_1$; $m_{z,1}(s_1)$ and $s_{z,1}^2(s_1)$ are the sample mean and variance of responses among subjects with baseline value S_1 who were assigned to z . Non-informative priors are assumed for normal and multinomial parameters, as well as ignorability given S_1 ; setting $B = \{F_1(s_1)\}$ in (4a)–(4b) establishes the results with some calculation. Note that the square-bracketed expression in (6b) is the predictive covariance of $F_1(s_1)$ and $F_1(s'_1)$. This component of variance occurs because the $\{F_1(s_1)\}$ are not known *a priori* and must be estimated.

The results for $K = 1$ were developed for treating non-response in sample surveys. They apply directly to SR if we consider adaptive randomized assignment as a missing data mechanism, which determines whether the potential Y_z response for any given subject is observed. In

contrast to the uncontrolled non-response mechanisms underlying survey data, the SR mechanism is a known function of S_1 , which leads to more efficient inference.

To make this idea precise, let E^* and V^* denote expectation and variance, respectively, conditional on the observed Y_z responses and S_1 values. Little and Rubin derive the terms of the summands in (6a) as follows:

$$\begin{aligned} E^*[F_1^2(s_1)s_{z,1}^2(s_1)/n_{z,1}(s_1)] &= [E^*(F_1(s_1))^2 + V^*(F_1(s_1))]s_{z,1}^2(s_1)/n_{z,1}(s_1) \\ &= [f_1^2(s_1) + n^{-1}f_1(s_1)\{1 - f_1(s_1)\}]s_{z,1}^2(s_1)/n_{z,1}(s_1) \end{aligned}$$

If E^* and V^* additionally include conditioning on the known randomization probabilities $\{\Gamma_{z,1}(s_1)\}$, where $\Gamma_{z,1}(s_1) = \text{pr}(z|S_1 = s_1)$, then

$$\begin{aligned} E^*[F_1^2(s_1)s_{z,1}^2(s_1)/n_{z,1}(s_1)] &= E^*[F_1^2(s_1)/\{\pi NF_1(s_1)\Gamma_{z,1}(s_1)\}]s_{z,1}^2(s_1) \\ &= E^*[F_1(s_1)/\Gamma_{z,1}(s_1)]s_{z,1}^2(s_1)/n \\ &= f_1(s_1)s_{z,1}^2(s_1)/\{n\Gamma_{z,1}(s_1)\} \end{aligned} \quad (7)$$

for a sample of size n ; $n_{z,1}(s_1) = \pi M_{z,1}(s_1)$ is the ignorability assumption underlying the predictive calculation for (6), where $M_{z,1}(s_1) = NF_1(s_1)\Gamma_{z,1}(s_1)$ (gives the first equality), and π is the probability of selection for the sample, which we assume to be n/N , ignoring variation of order $1/N$ (gives the second equality). The third equality requires that the randomization probabilities are known. Note that (7) is equal to $f_1^2(s_1)s_{z,1}^2(s_1)/n_{z,1}(s_1)$, up to sampling variation; the second summand in (6a) has been eliminated *via* the additional conditioning.

Case $K > 1$: Suppose Y_z has a stratified normal distribution across subgroups indexed by possible state histories $\underline{s}_K$. The joint distribution of $\underline{s}_K$ is factored sequentially in the usual way, with intermediate outcomes $S_{z,k+1,i}$ distributed (conditional on $\underline{s}_K$) as i.i.d. multinomial random variables. We assume that parameters are distinct across state histories for a given time period and *a priori* independent across time periods. As before, non-informative priors are specified for normal and multinomial parameters.

The posterior calculations for $K > 1$ first proceed like those of Little and Rubin, by setting $B = \{P_{z,K}(\underline{s}_K)\}$ in (4), with all expectations and variances conditional on observed history (prior treatments and responses) as required by the sequential ignorability underlying SR. The predictive mean of $\bar{Y}_z$ is a version of the G -computational algorithm [15]:

$$E^*(\bar{Y}_z) = \sum_{\underline{s}_K} E^*(P_{z,K}(\underline{s}_K))m_{z,K}(\underline{s}_K) = \sum_{\underline{s}_K} \varphi_{z,K}(\underline{s}_K)m_{z,K}(\underline{s}_K) \quad (8)$$

where $m_{z,K}(\underline{s}_K)$ is the sample mean of final responses among subjects sequentially randomized to z through K and having state values $\underline{s}_{z,K} = \underline{s}_K$,

$$\varphi_{z,K}(\underline{s}_K) = \prod_{j=1}^K f_{z,j}(\underline{s}_j) \quad (9)$$

and $f_{z,j}(s_j)$ is the sample (conditional) response rate for $S_{z,j} = s_j$ given randomization to z through $j-1$ and $\underline{s}_{z,j-1} = \underline{s}_{j-1}$; $f_{z,1}(\underline{s}_1) \equiv f_1(s_1)$ as above for $K=1$. Note that the sample size (number of multinomial trials) for $f_{z,j}(s_j)$ is the number of subjects randomized to strategy z through $j-1$ with $\underline{s}_{z,j-1} = \underline{s}_{j-1}$, which we denote as $n_{z,j-1}(\underline{s}_{j-1})$; the corresponding multinomial count is the number of the $n_{z,j-1}(\underline{s}_{j-1})$ subjects with $S_{z,j} = s_j$. Like the G -computational algorithm, (9) results from iterative application of (4a), sequentially setting $B = \{P_{z,k}(\underline{s}_k)\}$ for the repeated (posterior) expectations.

The predictive variance decomposes into

$$V^*(\bar{Y}_z) = V_n^*(\bar{Y}_z) + V_p^*(\bar{Y}_z) \quad (10)$$

where

$$V_n^*(\bar{Y}_z) = \sum_{\underline{s}_K} [E^*(P_{z,K}^2(\underline{s}_K))] s_{z,K}^2(\underline{s}_K) / n_{z,K}(\underline{s}_K) = \sum_{\underline{s}_K} \varphi_{z,K}(\underline{s}_K) s_{z,K}^2(\underline{s}_K) / \{n \Gamma_{z,K}(\underline{s}_K)\} \quad (11)$$

$$V_p^*(\bar{Y}_z) = \sum_{\underline{s}_K, \underline{s}'_K} m_{z,K}(\underline{s}_K) m_{z,K}(\underline{s}'_K) \text{CV}^*(P_{z,K}(\underline{s}_K), P_{z,K}(\underline{s}'_K)) \quad (12)$$

and $s_{z,K}^2(\underline{s}_K) / n_{z,K}(\underline{s}_K)$ is the sample variance of $m_{z,K}(\underline{s}_K)$. As for $K=1$, we condition on the cumulative randomization probabilities $\{\Gamma_{z,K}(\underline{s}_K)\}$ to gain efficiency from the knowledge of the missing data mechanism generated by SR; the cross-sectional calculation used to obtain (7) extends easily to yield the second equality in (11). We refer to $V_n^*(\bar{Y}_z)$ as the ‘naïve’ variance estimator because it ignores the uncertainty due to estimation of the unknown $\{F_{z,k}(\underline{s}_k)\}$, $k \geq 1$; $V_p^*(\bar{Y}_z)$ is the ‘penalty’ paid for their not being known *a priori*.

Causal inference for $\bar{Y}_c - \bar{Y}_d$ further requires that we derive the predictive covariance of $\bar{Y}_c$ and $\bar{Y}_d$. Let CV^* denote covariance conditional on observed values, which includes randomization probabilities. By definition,

$$\text{CV}^*(\bar{Y}_c, \bar{Y}_d) = \sum_{\underline{s}_K, \underline{s}'_K} \text{CV}^*[P_{c,K}(\underline{s}_K) \bar{Y}_c(\underline{s}_K), P_{d,K}(\underline{s}'_K) \bar{Y}_d(\underline{s}'_K)] \quad (13)$$

where the summation over $\underline{s}_K$ and $\underline{s}'_K$ includes all possible pairs implied by the rules for covariance of two sums. To express the RHS of (13) in terms of sample quantities, fix $s_K, \underline{s}'_K$. Let $\Delta(\underline{s}_K, \underline{s}'_K, c, d) = \min\{\delta(\underline{s}_K, c, d), \delta(\underline{s}'_K, c, d)\}$, where

$$\delta(\underline{s}, c, d) = \min_k \{A_{c,k} \neq A_{d,k} | \underline{s}_{c,K} = \underline{s}_{d,K} = \underline{s}\} \quad (14)$$

and $\delta(s_K, c, d) = K + 1$ if treatments are common throughout. Equation (14) formulates the divergence of treatments for two strategies as a function of observable state histories. Suppressing the arguments to Δ , we define

$$P_{\Delta}(s_{\Delta}) \equiv P_{c,\Delta}(s_{\Delta}) = P_{d,\Delta}(s_{\Delta}), \quad P_{\Delta}(s'_{\Delta}) \equiv P_{c,\Delta}(s'_{\Delta}) = P_{d,\Delta}(s'_{\Delta}) \quad (15)$$

where s_{Δ}, s'_{Δ} are the first Δ values of s_K, s'_K , respectively. The following result shows that the predictive covariance of two strategies quantifies the lack of independence arising from their shared use of common data, as parameterized by Δ . It is straightforward to derive the predictive covariance by setting $B = \{P_{\Delta}(s_{\Delta}), P_{\Delta}(s'_{\Delta})\}$ in (4b); see the Appendix.

Result 1—Assume for strategies c and d that (i) sequential ignorability holds and (ii) $\delta(s_K, c, d) \leq K$ for all state histories s_K . Then $CV^*(\bar{Y}_c, \bar{Y}_d)$ can be expressed as

$$\sum_{s_K, s'_K} m_{c,K}(s_K) m_{d,K}(s'_K) \prod_{k=\Delta+1}^K f_{c,k}(s_k) f_{d,k}(s'_k) CV^*(P_{\Delta}(s_{\Delta}), P_{\Delta}(s'_{\Delta})) \quad (16)$$

Condition (ii) serves to simplify the presentation of results, and can be assumed with no loss of generality for our purposes. When it fails to hold, $CV^*(\bar{Y}_c, \bar{Y}_d)$ is equal to (16) plus additional terms that account for the stratified (naïve) variance due to subgroups with identical treatments under either strategy:

$$\sum_{s_K} I(\delta(s_K, c, d) = K + 1) \varphi_{z,K}(s_K) s_{z,K}^2(s_K) / \{n \Gamma_{z,K}(s_K)\}, \quad z = c \text{ or } d \quad (17)$$

We note that the potential responses Y_c and Y_d are themselves also likely correlated, even though this correlation is unobservable when treatments are distinct. However, we have previously shown for $K = 1$ that the effect of such unobservable correlation on the predictive covariance of $\bar{Y}_c$ and $\bar{Y}_d$ is negligible if N is large relative to n (see Appendix C of [16]). Hence for most applications, only observable correlation due to shared common data will impact SCI; see the Appendix for detailed proof.

The predictive covariance of $P_{\Delta}(s_{\Delta})$ and $P_{\Delta}(s'_{\Delta})$, which encodes the dependence between strategies c and d due to shared data for fixed s_K and s'_K , is obtained sequentially according to the multinomial models for intermediate outcomes. Assuming that there are only two strategies under evaluation, simple expressions result for the case $s_{\Delta} \neq s'_{\Delta}$:

$$CV^*(P_{\Delta}(s_{\Delta}), P_{\Delta}(s'_{\Delta})) = -\varphi_{\Delta}(s_{\Delta}) \varphi_{\Delta}(s'_{\Delta}) / n \quad (18)$$

and $s_{\Delta} = s'_{\Delta}$:

$$CV^*(P_{\Delta}(s_{\Delta}), P_{\Delta}(s'_{\Delta})) = \varphi_{\Delta}(s_{\Delta}) (1 - \varphi_{\Delta}(s_{\Delta})) / n \quad (19)$$

where

$$\varphi_{\Delta}(\underline{s}_{\Delta}) = \prod_{j=1}^{\Delta} f_j(\underline{s}_j), \quad f_j(\underline{s}_j) \equiv f_{c,j}(\underline{s}_j) = f_{d,j}(\underline{s}_j), \quad j=1, \dots, \Delta$$

and n is the number of subjects in the trial. Underlying the derivation of (18)–(19) is the guarantee that subjects are never randomized prior to Δ . In general, the expression for $CV^*(P_{\Delta}(\underline{s}_{\Delta}), P_{\Delta}(\underline{s}'_{\Delta}))$ does not simplify, and is expressed in terms of a sequence of components of covariance, according to the inductive calculations in the Appendix. Analogously, $V_p^*(\bar{Y}_z)$ is obtained inductively (set $\Delta = K$), to account for the uncertainty due to unknown $\{F_{z,k}(\underline{s}_k)\}$, $k \geq 1$, when sequentially estimating the subgroup proportions $P_{z,K}(\underline{s}_K)$.

We note that the coverage properties of the predictive covariance can be obtained from those of the predictive variance, by converting $CV^*(\bar{Y}_c, \bar{Y}_d)$ to $V_p^*(\bar{Y}_z)$ for a suitably chosen strategy $z = z(c, d)$; see the Appendix. Simulation shows that the inductive calculations underlying $V_p^*(\bar{Y}_z)$ for general z adequately account for the variance due to sequential estimation of subgroup proportions [13].

4. COMPARISON TO OTHER APPROACHES FOR CAUSAL INFERENCE

4.1. SCI vs BI for pairwise effects

Previously, we showed that when SR is used to assign subjects to adaptive treatment strategies, the resulting data can be allocated to separate groups, one per strategy [12]. Such allocation permits standard BI to be used for valid estimation of causal effects. This naturally raises the question about the relative efficiency (RE) of the two inference approaches for evaluating adaptive treatment strategies. We consider here the variances under SCI and BI when subjects are randomized sequentially; both causal estimators are consistent for pairwise effects, with BI also producing an unbiased estimator. To simplify derivations, we assume that SR is carried out so that $\{\Gamma_{z,k}(\underline{s}_k)\} = \{\gamma_{z,k}(\underline{s}_k)\}$, $k \geq 1$, for $z = c, d$, where $\{\gamma_{z,k}(\underline{s}_k)\}$ are the observed randomization proportions corresponding to $\{\Gamma_{z,k}(\underline{s}_k)\}$. This condition might be obtained *via* sequentially blocked randomization.

It is possible to analytically relate the variances of the SCI and BI causal estimators when there are only two strategies under evaluation and they specify the same decision rules through $K - 1$. In this case, SCI is always at least as efficient as BI, with equality holding if $m_{z,K}(\underline{s}_K) \equiv m_z$, under each strategy $z = c, d$, i.e. there is no between-subgroup response heterogeneity; see the Appendix for the proof.

Result 2—Assume for strategies c and d that (i) sequential ignorability holds, (ii) $\delta(\underline{s}_K, c, d) = K$ is constant across state histories, and (iii) subjects are sequentially randomized either to $z = c$ or $z = d$ with probability p and $1 - p$, respectively. Then $V_{BI}(\bar{Y}_c - \bar{Y}_d) - V^*(\bar{Y}_c - \bar{Y}_d)$ is approximately:

$$\sum_{\underline{s}_K} n^{-1} \varphi_K(\underline{s}_K) [w_p \{m_{c,K}(\underline{s}_K) - E^*(\bar{Y}_c)\} + w_p^{-1} \{m_{d,K}(\underline{s}_K) - E^*(\bar{Y}_d)\}]^2 \quad (20)$$

where $V_{BI}(\bar{Y}_c - \bar{Y}_d)$ denotes the variance estimator under BI, and $w_p = \sqrt{\frac{1-p}{p}}$. The approximation arises because subgroup quantities under BI are only equal to those under SCI up to sampling variation, which we expect to be negligible under conditions (ii) and (iii). Result 2 shows that

efficiency gains due to shared use of values through K are governed by between-subgroup differences in Y_z , $z = c, d$.

To gauge the size of (20), we generated data under conditions (i)–(iii) and set $p = \frac{1}{2}$. Each simulation is replicated 2000 times with $K = 3$, $n = 500$, and $Y_z = \beta_{z,1}S_{z,1} + \beta_{z,2}S_{z,2} + \beta_{z,3}S_{z,3}$ plus (within-stratum) error that is distributed as $N(0, 4)$, $z = c, d$. The regression coefficients are $\beta_c = (0, 1, 2)$ and $\beta_d = (1, 2, 3)$, with higher Y_z values interpreted as better outcomes. Baseline states are equiprobable, with $S_1 = 1, 2$, or 3 ; $S_{z,2}$ and $S_{z,3}$ take on the same values and are produced according to the following transition matrix, which is consistent with ‘healthier’ subjects having a greater probability of better intermediate outcomes:

$$\begin{array}{ccc} 0.4 & 0.3 & 0.3 \\ 0.3 & 0.3 & 0.4 \\ 0.2 & 0.2 & 0.6 \end{array} \quad (21)$$

See [13] for more details on the simulation design.

The ‘true’ (across 2000 simulations) variances of $\bar{Y}_c - \bar{Y}_d$ under BI and SCI are 0.10 and 0.041, respectively; the means of $V_{BI}(\bar{Y}_c - \bar{Y}_d)$ and $V^*(\bar{Y}_c - \bar{Y}_d)$ across simulated data sets are 0.10 and 0.042. All predictive estimators exhibit modest bias, which mostly cancels out: mean of $V^*(\bar{Y}_c) = 0.027$, ‘true’ value = 0.026; mean of $V^*(\bar{Y}_d) = 0.044$; ‘true’ value = 0.044; mean of $CV^*(\bar{Y}_c, \bar{Y}_d) = 0.015$, ‘true’ value = 0.014. The gain in efficiency from SCI is substantial: mean RE = 0.51, where $RE = V^*(\bar{Y}_c - \bar{Y}_d)/V_{BI}(\bar{Y}_c - \bar{Y}_d)$.

Condition (ii) in Result 2 gives SCI the greatest advantage over BI because the predictive covariance increases, while within-strategy variances decrease, as the two strategies share more data. In contrast, the condition $\delta(\underline{s}_K, c, d) \equiv 1$ is most favorable to BI since S_1 is the only data common to both strategies when it holds. (Hereafter, we suppress the arguments $\{c, d\}$ to δ) To make this concrete, suppose strategies c' and d' are as described above for the simulated strategies c and d , but $\delta(\underline{s}_K) \equiv 1$, rather than $\delta(\underline{s}_K) \equiv 3$. As constructed, the distinct treatments of c' and d' produce the same intermediate outcomes; differential weighting of prior responses *via* the regression coefficients $\beta_{c'} = (0, 1, 2)$ and $\beta_{d'} = (1, 2, 3)$ generates the causal difference. Simulated results show the impact of $\delta(\underline{s}_K)$ on the precision of predictive causal estimates to be considerable: the mean RE (across 2000 simulations) is 0.97, approximately double that obtained under the simulation of (20).

4.2. SCI vs MM for pairwise effects

MM models offer a semi-parametric approach for evaluating adaptive treatment strategies *via* SR [2]. Under appropriate experimental conditions, SCI and MM provide the same sequential estimator of $\bar{Y}_c - \bar{Y}_d$, raising the issue of how their variance estimators compare (e.g. in terms of bias). The MM estimator is derived from an estimating formula as

$$\sum_{i=1}^n w_{c,i} Y_i - \sum_{i=1}^n w_{d,i} Y_i, \quad w_{z,i} = D_{z,i} \Gamma_{z,K}(\underline{s}_{K,i})^{-1} / \sum_{j=1}^n D_{z,j} \Gamma_{z,K}(\underline{s}_{K,j})^{-1}, \quad z=c, d \quad (22)$$

where the Y_i are the observed individual final responses; $D_z = 1$ indicates adherence to strategy z through K , and is otherwise 0. As above for BI, we assume for each strategy that $\{\Gamma_{z,k}(\underline{s}_k)\} = \{\gamma_{z,k}(\underline{s}_k)\}$, in which case it is straightforward to show that the MM causal estimator is equal to the difference in predictive means $E^*(\bar{Y}_c) - E^*(\bar{Y}_d)$.

The MM variance estimator for $\bar{Y}_c - \bar{Y}_d$, denoted $V_{MM}(\bar{Y}_c - \bar{Y}_d)$, is derived semi-parametrically-as

$$\frac{1}{n^2} \left\{ \sum_{i=1}^n D_{c,i} [\Gamma_{c,K}(\underline{s}_{K,i})^{-1} (Y_i - m_c)]^2 + \sum_{i=1}^n D_{d,i} [\Gamma_{d,K}(\underline{s}_{K,i})^{-1} (Y_i - m_d)]^2 - 2 \sum_{i=1}^n D_{c,i} D_{d,i} [\Gamma_{c,K}(\underline{s}_{K,i})^{-1} (Y_i - m_c)] [\Gamma_{d,K}(\underline{s}_{K,i})^{-1} (Y_i - m_d)] \right\} \quad (23)$$

where $m_z = E^*(\bar{Y}_z)$, $z = c, d$. Note that the cross-product terms account for dependence in pairwise comparisons due to subjects that receive the same treatments under either strategy [2]; if there are no subjects whose data are consistent with both c and d , (23) reduces to the sum of the MM within-strategy variances.

The MM estimator is calculated as a simple summation, and, similar to ANOVA estimators, $V_{MM}(\bar{Y}_c - \bar{Y}_d)$ decomposes into within- and between-subgroup components of variance. In particular, $V_{MM}(\bar{Y}_c - \bar{Y}_d)$ is minimized when there is no between-subgroup response heterogeneity under either strategy, i.e. $m_{z,K}(\underline{s}_K) \equiv m_z$, $z = c, d$. Assuming $\delta(\underline{s}_K) \leq K$ for all state histories $\underline{s}_K$, the MM estimator reduces in this case to

$$V_n^*(\bar{Y}_c) + V_n^*(\bar{Y}_d)$$

The analogous condition for $V^*(\bar{Y}_c - \bar{Y}_d)$ is that $\{F_{z,k}(\underline{s}_k)\}$, $k \geq 1$, must be known, in which case $V^*(\bar{Y}_z) = V_n^*(\bar{Y}_z)$ for each strategy z . These extreme cases illustrate the primary distinction between MM and SCI. The predictive variance estimator is calculated inductively as a sequence of components of variance and covariance to account for estimation uncertainty due to unknown $\{F_{z,k}(\underline{s}_k)\}$; the semi-parametric variance estimator is derived *via* a marginal model in a single step to account for (weighted) deviations from the grand mean.

The distinct approaches of SCI and MM bear on their ability to handle shared data. The sequential approach of SCI permits it to accommodate partially overlapping data between two strategies, which is not possible with the MM cross-product terms. Not fully accounting for dependence among intermediate outcomes can lead to overstated uncertainty for $\bar{Y}_c - \bar{Y}_d$. Furthermore, the dissimilar conditions under which $V_{MM}(\bar{Y}_c - \bar{Y}_d)$ and $V^*(\bar{Y}_c - \bar{Y}_d)$ are minimized hold true for within-strategy variances, indicating the potential for MM to produce quite different, possibly biased standard errors for $E^*(\bar{Y}_z)$. The bias properties of the MM variance estimators can be examined analytically and empirically through a connection to BI. Specifically, given $\delta(\underline{s}_K) \leq K$ for every observable state history, MM and BI produce the same result (up to negligible sampling variation) when there are only two strategies under evaluation. This is true because all subjects are randomized once under SR at some decision point to one of the two strategies, and MM treats adherence to a particular strategy non-sequentially *via* the indicator D_z . This effectively partitions the data into strategy-specific groups according to full compliance, using the same allocation probabilities as baseline randomization. Calculations for (22) and (23) then reduce to those for BI.

Consider the case when $\delta(\underline{s}_K) \equiv K$. Result 2 translates directly to the comparison of the MM variance estimator with the true variance by replacing sample quantities with their population counterparts, the (modified) difference (20) now representing the bias in $V_{MM}(\bar{Y}_c - \bar{Y}_d)$. This follows because the true variance of the predictive mean in repeated sampling is derived analogously to $V^*(\bar{Y}_z)$; the same holds true for the true covariance with respect to the predictive

covariance. Specifically, the calculation for the true variance for strategy z , denoted $V(E^*(\bar{Y}_z)|\mathbf{Y}_z, \underline{\mathbf{S}}_{z,K})$, sets B in (4b) to $\{n_{z,k}(\underline{s}_k)\}$, and the vector of corresponding multinomial counts, where $\mathbf{Y}_z, \underline{\mathbf{S}}_{z,K}$ are the complete population responses. (We assume n is large enough that $\{n_{z,k}(\underline{s}_k)\}$ are positive across repeated samples.) The substantial efficiency gains of SCI relative to BI demonstrated in the accompanying simulation example for (20) can also be re-interpreted as substantial positive bias in the MM estimator. Simple calculation shows that the bias in this example is not solely due to omitted covariance terms but also due to overestimated variances for the strategies. Because the MM within-strategy variance estimators, denoted $V_{\text{MM}}(\bar{Y}_z)$, $z = c, d$, are empirically the same as the BI ones, as we expect them to be, we do not attribute this component of bias to small inverse weights, sometimes responsible for making MM-type approaches high-variance estimators [14]. Rather, the bias reflects the marginal ‘one-step’ nature of the MM approach, which is unable to adequately capture the uncertainty underlying the sequential identification of the population mean. The modified (translated) version of Result 2 shows that the lack of between-subgroup response heterogeneity in the population eliminates this bias.

When the two strategies under evaluation lead to a single (sequential) randomization for each state history $\underline{s}_K$, we can further exemplify between-subgroup heterogeneity as a determinant of bias. As indicated above, MM and standard baseline calculations are the same under this scenario; $V_{\text{MM}}(\bar{Y}_c - \bar{Y}_d)$ also reduces to the sum of MM within-strategy variances. Moreover, because the subgroup proportions $\{P_{z,K}(\underline{s}_K)\}$, $z = c, d$, are identified non-sequentially when there are separate baseline groups, BI is formally equivalent to the cross-sectional ($K = 1$) case described in Section 3.1. This allows us to apply rule (4b) non-iteratively, setting $B = \{P_{z,K}(\underline{s}_K)\}$ to decompose each within-strategy variance under BI (and correspondingly under MM) into naïve and penalty terms, with direct comparison to SCI, or the true variance (after taking expectations). Little and Rubin [14] showed that the penalty term (6b) for $K = 1$ can be re-expressed in terms of between-subgroup variance. The interpretation of their result in our context generalizes Result 2, which uses response heterogeneity to more narrowly characterize the relative efficiency of BI to SCI (sample version), or the bias in MM variances (population version). In particular, given $\delta(\underline{s}_K) \leq K$ for every observable state history, lack of between-subgroup variance in the sample for both strategies implies that BI (and MM) is as or more efficient for pairwise effects than SCI, and in the population implies that the bias in $V_{\text{MM}}(\bar{Y}_z)$, $z = c, d$, is non-positive. These relationships hold only approximately under realistic scenarios of small response heterogeneity.

We remark that when population subgroup means under a strategy z are equal, the optimal (most efficient) estimating formula, defined in the Appendix of [2], reduces to that used to obtain the MM estimator of $\bar{Y}_z$ in (22); hence, the optimal variance estimator will be negatively biased when this condition holds if $V_{\text{MM}}(\bar{Y}_z)$ is.

5. RANDOMIZATION TO MULTIPLE ADAPTIVE STRATEGIES

We have considered in detail, both analytically and empirically, the case when there are only two strategies under evaluation. Although researchers may limit themselves to such studies, the principles of optimal sequential decision making dictate that ‘all possible’ strategies be studied simultaneously [2,17]. The notion of all possible strategies is a mathematical one and usually clinical principles will dictate which strategies are relevant for study [1]. The set of relevant strategies can typically be represented sequentially as a decision tree. Whether or not resources permit exhaustive investigation of all ‘branches’ of the tree, it may often be the case that multiple strategies will be included for study. We extend the discussion of Section 4 to this case.

5.1. SCI vs BI for pairwise effects—randomization to multiple strategies

When subjects are sequentially randomized to multiple adaptive treatment strategies, it is straightforward to see how this affects SCI precision for pairwise comparisons. Specifically, the sample sizes for sequentially derived predictive estimators are determined according to $\{\Gamma_{z,k}(\underline{S}_k)\}$, which reflect the number of treatment options provided at each decision point by the strategies under evaluation. Hence, the SR data available for comparing two strategies may be ‘censored’ as subjects are randomized to other strategies with previously common decision rules. In this way, the number of randomizations largely governs the loss of SCI precision due to inclusion of additional strategies in the study. In contrast, the loss of BI precision depends in the usual way on the number of strategies, with allocation of SR data to strategy-specific groups determined by the randomization probabilities. Because the number of strategies may be much larger than the number of randomizations faced by any given subject in the trial, there is potential for SCI to be much more efficient than BI. We illustrate this next with three related scenarios.

Assume $K = 2$ and suppose first that subjects are adaptively treated according to the decision rules for trinary states in Table I, where $\{SW_i\}$ and $\{AUG_i\}$ are the options for switching and augmenting existing treatment, respectively. Each strategy specifies an initial switch option and an augmentation option, giving rise to four different possibilities. As Table I shows, many patients are randomized either once or not at all. This not only reduces sequential ‘censoring’ due to the inclusion of multiple treatment strategies in the study, but also increases the SR data shared by different strategies. Both factors should produce substantial gains in the RE of BI to SCI; simulation bears this out. Given the setup in Section 4.1 with equiprobable baseline states and transition matrix (21), we specify the regression of Y_z on $\underline{S}_{z,K}$, $z = c, d$, via $\beta_c = (0, 1)$, $\beta_d = (1, 2)$. The mean RE for comparing (SW_1, AUG_1) with (SW_2, AUG_2) , or analogously for comparing (SW_1, AUG_2) with (SW_2, AUG_1) , is 0.49. Comparison of strategies with common switch or augmentation options would be even more favorable to SCI since subjects would be randomized at most once.

Assume next that S_k is binary, rather than trinary, and that rows 1–2, 4–5 for trinary states in Table I specify the decision rules. As before, there are four possible strategies, but now all subgroups but one are randomized twice. In this case, the mean RE value is 0.92, given the transition matrix

$$\begin{pmatrix} 0.4 & 0.6 \\ 0.6 & 0.4 \end{pmatrix}$$

aggregated from (21). We can further modify this scenario so that there are two treatment options for all subgroups at each stage, thereby ensuring that every subject is randomized twice. The second subtable in Table I depicts a version of this structure with four possible strategies, parameterized by the choice of initial medication (M_1 or M_2) and follow-up therapy (T_1 or T_2). There are no shared data among strategies other than S_1 . Nonetheless, the mean RE is 0.93, only a modest gain in efficiency for BI.

In general, the number of randomizations under SR is limited by the number of decision points K , whereas the number of possible strategies can grow multiplicatively with K . However, because adaptive strategies are clinical algorithms, the number of possible strategies is not determined as a combinatoric function of K and can vary widely for any given number of decision points. Consider again the case of binary S_k with $K = 2$, but assume that there are distinct pairs of treatment options to which subjects are (adaptively) randomized for each value of S_1 and $(\underline{S}_2, A1)$, all possible combinations of interest. Then the number of strategies is 64,

rather than the four strategies portrayed in the third example. Moreover, the sample sizes for predictive quantities are still determined by two randomizations per subject; as a result, the sample sizes are the same whether there are 64 or four strategies. For pairwise comparisons with distinct initial options, BI is approximately $\frac{1}{16}$ as efficient as SCI when there are 64 strategies, rather than having RE = 0.93, as in the previous example. For larger values of K , the efficiency gains with SCI can be even more striking.

5.2. SCI vs MM for pairwise effects—randomization to multiple strategies

The connection between BI and MM inference derived in Section 4.2 holds when multiple strategies are under evaluation, but requires that the final SR probabilities are the same across state histories for each strategy z , i.e. $\{\Gamma_{z,K}(\underline{s}_K)\} \equiv \Gamma_{z,K}$, with $\Sigma \Gamma_{z,K} = 1$. Analogous to the two-strategy case, this implies that the bias in $V_{MM}(\bar{Y}_z)$ will be non-positive whenever there is no response heterogeneity in the population under strategy z . To compare SCI and MM approaches under more general scenarios, we proceed empirically and assume that subjects are randomized with equal probability to one of two treatment options at each decision point. We vary two features of the simulation setup in Section 4.1 to examine the performance of the two variance estimators. First, we set $K = 2$ and 4, as larger values of K may increase uncertainty in predictive variances arising from sequential estimation of the subgroup proportions $P_{z,K}(\underline{s}_K)$. The regression coefficients for Y_z , $z = c, d$, are: $\beta_c = (0, 1)$, $\beta_d = (1, 2)$ for $K = 2$; $\beta_c = (0, 1, 2, 3)$, $\beta_d = (1, 2, 3, 4)$ for $K = 4$. Second, we consider two transition matrices to vary between-subgroup response heterogeneity, which affects the bias in $V_{MM}(\bar{Y}_z)$, $z = c, d$. (Note that the distribution of $\underline{s}_{z,K}$, specified *via* the transition matrix, governs the between-subgroup variance for fixed β_z .) We use

$$\begin{array}{ccc} 0.60 & 0.20 & 0.20 \\ 0.10 & 0.60 & 0.30 \\ 0.01 & 0.09 & 0.90 \end{array} \quad (24)$$

which increases stratified response heterogeneity over time, and

$$\begin{array}{ccc} 0.35 & 0.35 & 0.30 \\ 0.35 & 0.35 & 0.30 \\ 0.30 & 0.35 & 0.35 \end{array} \quad (25)$$

which promotes mixing across subgroups, thereby reducing between-subgroup variance over time. As in the original setup, we keep the within-stratum (-subgroup) variance of Y fixed.

The results in Table II show that the MM variance estimator may be substantially more biased and more variable than SCI. The degree of response heterogeneity, reflected in the choice of transition matrix, is an important determinant of this bias. Rows 3 and 4 of the first subtable show MM bias to be quite pronounced for larger K and β_z , e.g. relative bias exceeds 180% when $K = 4$, $\beta_z = (1, 2, 3, 4)$. By contrast, if subjects are sequentially randomized once to one of two strategies (as in Section 4) and $\delta(\underline{s}_K) \equiv 1$, then $V_{MM}(\bar{Y}_z) = 0.18$ (true value = 0.14; standard error = 0.014) when $K = 4$, $\beta_z = (1, 2, 3, 4)$ and (24) is used; the relative bias is considerably less, even though it exceeds 20%. Note that relative bias is partly a function of the true naïve variance, which depends on the choice of within-stratum variance (set to 4 here).

6. DISCUSSION

In this paper, we have formally unlinked the mode of inference for the comparative effects of adaptive treatment strategies from the mode of randomization used to obtain data. There is a natural appeal in adopting sequential approaches for both aspects of clinical trial methodology, since the interventions themselves develop sequentially over time. However, there are also challenges to such approaches. In particular, when adaptively determined treatments are assigned sequentially, each decision point that represents a choice between two (or more) strategies constitutes a point of randomization. This gives rise to a nested structure of treatment assignment over time, which in turn gives rise to a monotone structure of shared data among strategies. To accommodate the resulting statistical dependence, we have applied predictive (Bayesian) inference to estimate the covariance between two strategies. Similar to its variance counterparts, the predictive covariance is a sequential estimator, and can thus explicitly account for partially (or completely) overlapping SR data among strategies.

Standard (baseline) methods for comparing treatments offer a simpler alternative to SCI. The simulations reported in Section 4 for randomization to two strategies show that the loss in efficiency with BI relative to SCI could be substantial when the strategies have considerable overlap. When the strategies immediately specify distinct treatments for all baseline values, the RE of the two approaches depends on the distribution of intermediate and final responses, which is not known *a priori*. The results of Section 4.2 suggest that SCI is a sensible choice in these cases, because the between-subgroup variance will likely be positive, which in turn restricts BI efficiency. When subjects are randomized to multiple strategies, the efficiency gains with SCI can be sizeable, even for two decision points, largely due to the impact of shared SR data on within-strategy standard errors. We emphasize though that the degree of BI inefficiency depends intrinsically on the sequential structure of alternative decision rules, which in turn is determined more by clinical principles that dictate how the strategies are parameterized in terms of treatment options, than by combinatorics.

Semi-parametric alternatives to SCI, such as MM, provide a convenient marginal ‘one-step’ estimator for comparing strategies, but, by their very nature, are unable to accommodate dependence due to partial overlapping data among strategies. Moreover, the MM within-strategy variance estimator may fail to adequately capture the uncertainty underlying the sequential identification of the population mean. Sections 4 and 5 showed that causal inference *via* MM could suffer from overestimated or underestimated variances, influenced in a large part by response heterogeneity across subgroups defined by state histories. The simulation studies showed the bias in MM within-strategy variances to be quite modest when the between-subgroup variance was substantially reduced, yet this case corresponds to that when adjustment *via* sequential ignorability yields little bias reduction in population mean estimates. Furthermore, it is quite likely in designed randomized studies that the final outcome would be strongly related to intermediate ones, and that strategies would not be uniformly successful (or unsuccessful) across subgroups. The MM approach performs less well as the number of randomizations increases, in terms of both bias and variability of the variance estimator, if response heterogeneity increases with the number of decision points (as might be expected when adaptive treatment options under evaluation are considered to be in clinical equipoise). This is in direct contrast to SCI, which explicitly accounts for the sequential nature of SR data. For this reason, SCI will be at least as efficient as MM, unless response heterogeneity is small.

We note that it is straightforward to show that the data obtained *via* baseline randomization can be analyzed *sequentially* as though allocated *via* SR. Thus, the intricacies of the design do not necessarily impinge on the choice of the inferential method. In particular, SCI offers an attractive approach for causal inference, whether subjects are randomized to strategies at baseline or sequentially. The merits of either approach to randomization when treatments are

adaptive continue to be discussed and will likely depend on pragmatic as well as statistical issues [1–3,18].

This paper has focused exclusively on randomized trials, but carries considerable import for observational investigations as well. Many researchers still use non-randomized data for estimating the comparative effects of adaptive treatment strategies. When treatment selection is not controlled, the SCI and MM approaches produce the same estimator of causal effects. Hence, the results in Sections 4 and 5 apply, and are of particular note for sequentially assigned (observational) treatments. Previously, we showed that estimates of the population mean derived sequentially *via* the *G*-computational algorithm may also suffer from finite-sample bias [16], a result that has been studied recently through in-depth simulation [13]. Such bias applies to both SCI and MM estimators, and is of greatest concern when treatment choices are uncontrolled and vary widely over time. This paper reinforces the message that the validity of causal inference from observational data may rest as much with the properties of the estimator as with the plausibility of the assumption that there are no omitted (unobserved) confounders.

Acknowledgments

This work was supported by the National Institute of Mental Health Grant No. R01-MH51481 to Stanford University.

Contract/grant sponsor: National Institute of Mental Health; contract/grant number: R01-MH51481

References

1. Lavori PW, Dawson R. Dynamic treatment regimes: practical design considerations. *Clinical Trials* 2004;1:9–20. [PubMed: 16281458]
2. Murphy S. An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine* 2005;24:1455–1481. [PubMed: 15586395]
3. Thall PF, Millikan RE, Sung HG. Evaluating multiple treatment courses in clinical trials. *Statistics in Medicine* 2000;19:1011–1028. [PubMed: 10790677]
4. Schneider LS, Tariot PN, Lyketsos CG, Dagerman KS, Davis KL, Davis S, Hsiao JK, Jeste DV, Katz IR, Olin JT, Pollock BG, Rabins PV, Rosenheck RA, Small GW, Lebowitz B, Lieberman JA. National Institute of Mental Health Clinical Antipsychotic Trials of Intervention Effectiveness (CATIE): Alzheimer disease trial methodology. *American Journal of Geriatric Psychiatry* 2001;9:346–360. [PubMed: 11739062]
5. Brooner RK, Kidorf M. Using behavioral reinforcement to improve methadone treatment participation. *Science and Practice Perspectives* 2002;1:38–48. [PubMed: 18567965]
6. Rush A, Fava M, Wisniewski S, et al. Sequenced Treatment Alternatives to Relieve Depression (STAR*D): rationale and design. *Controlled Clinical Trials* 2004;25(1):119–142. [PubMed: 15061154]
7. Lieberman JA, Stroup TS, McEvoy JP, Swartz MS, Rosenheck RA, Perkins DO, Keefe RS, Davis SM, Davis CE, Lebowitz BD, Severe J, Hsiao JK. Clinical Antipsychotic Trials of Intervention Effectiveness (CATIE) Investigators: effectiveness of antipsychotic drugs in patients with chronic schizophrenia. *New England Journal of Medicine* 2005;353:1286–1288. [PubMed: 16172204]
8. Lavori PW, Dawson R, Rush AJ. Flexible treatment strategies in chronic disease: clinical and research implications. *Biological Psychiatry* 2000;48:605–614. [PubMed: 11018231]
9. Lavori PW, Rush JA, Wisniewski S, et al. Strengthening clinical effectiveness trials: equipoise-stratified randomization. *Biological Psychiatry* 2001;50:792–801. [PubMed: 11720698]
10. Dawson R, Lavori PW. Placebo-free designs for evaluating new mental health treatments: the use of adaptive treatment strategies. *Statistics in Medicine* 2004;23:3249–3262. [PubMed: 15490427]
11. Dawson R, Lavori PW, Luby J, Ryan N, Geller BG. Adaptive strategies for treating childhood mania. *Biological Psychiatry* 2007;61:758–764. [PubMed: 17166485]
12. Dawson R, Lavori PW. Comparison of designs for adaptive treatment strategies: classical vs. adaptive randomization. *Journal of Statistical Planning & Inference* 2003;117:365–385.

13. Lavori PW, Dawson R. Improving the efficiency of estimation in randomized trials of adaptive treatment strategies. *Clinical Trials*. 2007;in press
14. Little, RJA.; Rubin, DB. *Statistical Analysis with Missing Data*. Wiley; New York: 1987.
15. Robins, JM. The analysis of randomized and nonrandomized AIDS treatment trials using a new approach to causal inference in longitudinal studies. In: Sechrest, L.; Freeman, H.; Mulley, A., editors. *Health Service Research Methodology: A Focus on AIDS*. NCHSR, U.S. Public Health Service; Washington, DC, U.S.A: 1989. p. 113-159.
16. Dawson R, Lavori PW. Using inverse weighting and predictive inference to estimate the effects of time-varying treatments on the discrete-time hazard. *Statistics in Medicine* 2002;21:1641–1661. [PubMed: 12111903]
17. Lavori PW, Dawson R. A design for testing clinical strategies: biased-coin adaptive within-subject randomization. *Journal of the Royal Statistical Society, Series A* 2000;163:29–38.
18. Greenhouse JB, Meyer MM. A note on randomization and selection bias in maintenance therapy clinical trials. *Psychopharmacology Bulletin* 1991;27:225–229. [PubMed: 1775592]

Appendix

APPENDIX A

Proofs and derivations

Proof of Result 1

For the proof, fix $\underline{s}_K, \underline{s}'_K$ throughout, and apply (4b) by setting $B=\mathbf{P}_\Delta=\{P_\Delta(\underline{s}_\Delta), P_\Delta(\underline{s}'_\Delta)\}$. Then

$$\begin{aligned} & CV^*\{P_{c,K}(\underline{s}_K)\bar{Y}_c(\underline{s}_K), P_{d,K}(\underline{s}'_K)\bar{Y}_d(\underline{s}'_K)\} \\ &= E^*[CV^*\{P_{c,K}(\underline{s}_K)\bar{Y}_c(\underline{s}_K), P_{d,K}(\underline{s}'_K)\bar{Y}_d(\underline{s}'_K)|\mathbf{P}_\Delta\}] \end{aligned} \quad (\text{A1})$$

$$+ CV^*[E^*\{P_{c,K}(\underline{s}_K)\bar{Y}_c(\underline{s}_K)|\mathbf{P}_\Delta\}, E^*\{P_{d,K}(\underline{s}'_K)\bar{Y}_d(\underline{s}'_K)|\mathbf{P}_\Delta\}] \quad (\text{A2})$$

$$= E^*[P_\Delta(\underline{s}_\Delta)P_\Delta(\underline{s}'_\Delta)]CV^*\{P_{c,K-\Delta}(\underline{s}_{K-\Delta})\bar{Y}_c(\underline{s}_K), P_{d,K-\Delta}(\underline{s}'_{K-\Delta})\bar{Y}_d(\underline{s}'_K)\} \quad (\text{A3})$$

$$+ E^*[P_{c,K-\Delta}(\underline{s}_{K-\Delta})\bar{Y}_c(\underline{s}_K)]E^*[P_{d,K-\Delta}(\underline{s}'_{K-\Delta})\bar{Y}_d(\underline{s}'_K)]CV^*\{P_\Delta(\underline{s}_\Delta), P_\Delta(\underline{s}'_\Delta)\} \quad (\text{A4})$$

where $P_{c,K-\Delta}(\underline{s}_{K-\Delta}) = P_{c,K}(\underline{s}_K)/P_\Delta(\underline{s}_\Delta)$; $P_{d,K-\Delta}(\underline{s}'_{K-\Delta})$ is defined similarly. We claim that (A3) is zero or approximately so, assuming negligible finite population corrections. This follows from the predictive covariance of ‘post-treatment’ potential outcomes being zero (see below), where post-treatment here means ‘post’ Δ . Thus, the component of predictive covariance due to $\underline{s}_K, \underline{s}'_K$ is equal to (A4), which becomes

$$m_{c,K}(\underline{s}_K)m_{d,K}(\underline{s}'_K) \prod_{k=\Delta+1}^K f_{c,k}(\underline{s}_k)f_{d,k}(\underline{s}'_k)CV^*(P_\Delta(\underline{s}_\Delta), P_\Delta(\underline{s}'_\Delta)) \quad (\text{A5})$$

given the models specified in Section 3, thereby establishing the result. When condition (ii) for Result 1 fails to hold, the predictive covariance of $\bar{Y}_{c,K}(\underline{s}_K)$ and $\bar{Y}_{d,K}(\underline{s}'_K)$ reduces to $I(\underline{s}_K = \underline{s}'_K) V^*(\bar{Y}_{c,K}(\underline{s}_K))$. The additional terms in (17) account for the ‘naïve’ components due to identical treatments under either strategy.

Proof that (A3) is approximately zero

Fix $\underline{s}_K, \underline{s}'_K$ as for (A1); we show that

$$CV^*\{P_{c,K-\Delta}(\underline{s}_{K-\Delta})\bar{Y}_c(\underline{s}_K), P_{d,K-\Delta}(\underline{s}'_{K-\Delta})\bar{Y}_d(\underline{s}'_K)\} \rightarrow 0$$

as $N_{z,k}(s)/n_{z,k}(s) \rightarrow \infty$, for $k > \Delta$, and $z = c, s = \underline{s}_K; z = d, s = \underline{s}'_K$ to establish the result. First, set $B = \mathbf{P}_{K-\Delta} = \{P_{c,K-\Delta}(\underline{s}_{K-\Delta}), P_{d,K-\Delta}(\underline{s}'_{K-\Delta})\}$ in (4b) to obtain

$$\begin{aligned} & CV^*\{P_{c,K-\Delta}(\underline{s}_{K-\Delta})\bar{Y}_c(\underline{s}_K), P_{d,K-\Delta}(\underline{s}'_{K-\Delta})\bar{Y}_d(\underline{s}'_K)\} \\ &= E^*[P_{c,K-\Delta}(\underline{s}_{K-\Delta}), P_{d,K-\Delta}(\underline{s}'_{K-\Delta})] CV^*\{\bar{Y}_c(\underline{s}_K), \bar{Y}_d(\underline{s}'_K)\} \end{aligned} \quad (A6)$$

$$+ E^*[\bar{Y}_c(\underline{s}_K)] E^*[\bar{Y}_d(\underline{s}'_K)] CV^*\{P_{c,K-\Delta}(\underline{s}_{K-\Delta}), P_{d,K-\Delta}(\underline{s}'_{K-\Delta})\} \quad (A7)$$

The predictive covariance of $\bar{Y}_c(\underline{s}_K)$ and $\bar{Y}_d(\underline{s}'_K)$ in (A6) is well defined for subjects in the population such that $\underline{s}_{c,K} = \underline{s}_K$ and $\underline{s}_{d,K} = \underline{s}'_K$. (See [12] for a similar stratification of subjects by potential outcomes.) We specify the joint distribution of potential responses (Y_c, Y_d) for these subjects to be bivariate normal; the proof extends easily to discrete distributions. Consider first the case when $\underline{s}_K \neq \underline{s}'_K$. Assuming that the mean and covariance parameters are *a priori* independent, then

$$CV^*\{\bar{Y}_c(\underline{s}_K), \bar{Y}_d(\underline{s}'_K)\} \rightarrow 0 \quad (A8)$$

as $N_{z,k}(s)/n_{z,k}(s) \rightarrow \infty, z = c, s = \underline{s}_K; z = d, s = \underline{s}'_K$. (For proof, see Appendix C in [16], noting that given sequential ignorability, the conditions on $N_{z,k}(s)/n_{z,k}(s), z = c, s = \underline{s}_K; z = d, s = \underline{s}'_K$, are inherited by subjects for whom the predictive covariance is well defined.) The proof for $\underline{s}_K = \underline{s}'_K$ is analogous if $\Delta \leq K$, as required by condition (ii) for Result 1. The predictive covariance in (A7) is handled similarly, although it is necessary to derive the result sequentially through an inductive argument. The inductive step conditions on $\mathbf{P}_{\Delta+j} = \{P_{c,\Delta+j}(\underline{s}_{\Delta+j}), P_{d,\Delta+j}(\underline{s}'_{\Delta+j})\}$ to obtain

$$CV^*\{P_{c,\Delta+j+1}(\underline{s}_{\Delta+j+1}), P_{d,\Delta+j+1}(\underline{s}'_{\Delta+j+1})\} \rightarrow 0$$

as $N_{z,\Delta+j+1}(s)/n_{z,\Delta+j+1}(s) \rightarrow \infty$, $z = c$, $s = \underline{s}_{\Delta+j+1}$; $z = d$, $s = \underline{s}'_{\Delta+j+1}$, given that the limit holds for the predictive covariance of $\mathbf{P}_{\Delta+j}$.

Derivation of predictive covariance of $CV^*(P_{\Delta}(\underline{s}_{\Delta}), P_{\Delta}(\underline{s}'_{\Delta}))$

Fix $\underline{s}_K, \underline{s}'_K$ and proceed by induction on k . The case $k = 1$ is the bracketed expression in the predictive covariance derived for $K = 1$. The inductive step sets $B = \mathbf{P}_k = \{P_k(\underline{s}_k), P_k(\underline{s}'_k)\}$ in (4b), where $k < \Delta$ and strategy indicators are suppressed as in $\mathbf{P}_{\Delta}$, to obtain

$$\begin{aligned} & CV^*\{P_{k+1}(\underline{s}_{k+1}), P_{k+1}(\underline{s}'_{k+1})\} \\ &= E^*[P_k(\underline{s}_k), P_k(\underline{s}'_k)] CV^*\{F_{k+1}(\underline{s}_{k+1}), F_{k+1}(\underline{s}'_{k+1})\} \\ &+ f_{k+1}(\underline{s}_{k+1})f_{k+1}(\underline{s}'_{k+1}) CV^*\{P_k(\underline{s}_k), P_k(\underline{s}'_k)\} \end{aligned} \quad (A9)$$

noting that $P_{k+1}(\underline{s}_{k+1}) = F_{k+1}(\underline{s}_{k+1})P_k(\underline{s}_k)$. Assuming that CV^* includes conditioning on the known randomization probabilities, (A9) becomes

$$[I(s_{k+1} = \underline{s}'_{k+1})f_{k+1}(s_{k+1}) - f_{k+1}(s_{k+1})f_{k+1}(\underline{s}'_{k+1})]\varphi_k(\underline{s}_k)/\{n\Gamma_k(\underline{s}_k)\}$$

for $\underline{s}_k = \underline{s}'_k$ and is otherwise 0 by an *a priori* independence assumption, which completes the induction step; $\varphi_k(\underline{s}_k)$ is defined as in (9), setting $K = k$. The derivation applies to

$CV^*(P_{zK}(\underline{s}_K), P_{zK}(\underline{s}'_K))$ in formula (12) for $V_p^*(\bar{Y}_z)$, providing a formal justification for the inductive probability calculation of within-strategy variance in [13].

Derivation of the predictive covariance as a component of variance

To convert the predictive covariance for strategies c and d to a variance component, we define strategy c/d for each subject as follows:

$$\text{if } \underline{A}_{c,k,i} = \underline{A}_{d,k,i}, \text{ set } S_{c/d,k,i} = s_{c,k}, \quad A_{c/d,k,i} = A_{c,k,i} \quad (A10a)$$

$$\text{otherwise set } S_{c/d,k,i} = (s_{z,k}, z), \quad A_{c/d,k,i} = A_{z,k,i} \text{ for } S_{c/d,k,i} = (s_{z,k}, z), \quad z = c, d \quad (A10b)$$

where $S_{z,k,i} = s_{z,k}$, $z = c, d$. The notation $(s_{z,k}, z)$ re-labels output values in order to provide distinct multinomial values for $S_{c/d,k}$ and to ensure that treatments under the two strategies can be correspondingly mapped into adaptively determined treatments under c/d . Re-labeling intermediate outcome values in this fashion makes it automatic that the (conditional) multinomial distributions of $\{S_{c/d,k}\}$ inherit the randomization probabilities used to assign subjects to strategies. The final outcomes under c/d are similarly defined. The 'penalty' term of the predictive variance of c/d then coincides with the predictive covariance of c and d .

Proof of Result 2

Given conditions (i)–(iii) required for Result 2, it follows with some calculation that $V^*(\bar{Y}_c - \bar{Y}_d)$ decomposes into $V_n^*(\bar{Y}_c) + V_n^*(\bar{Y}_d)$ plus

$$1/n \left[\sum_{\underline{s}_K} \varphi_K(\underline{s}_K) \{m_{c,K}(\underline{s}_K) - m_{d,K}(\underline{s}_K)\}^2 - E^*(\bar{Y}_c - \bar{Y}_d)^2 \right] \quad (\text{A11})$$

where n is the number of subjects in the trial. To prove Result 2, we apply (18), (19) and (11) to re-express $V^*(\bar{Y}_c - \bar{Y}_d) - V_n^*(\bar{Y}_c) - V_n^*(\bar{Y}_d)$ as

$$\begin{aligned} & 1/n \left[\sum_{\underline{s}_K} \varphi_K(\underline{s}_K) \{1 - \varphi_K(\underline{s}_K)\} \{m_{c,K}(\underline{s}_K)^2 + m_{d,K}(\underline{s}_K)^2\} \right. \\ & - 2 \sum_{\underline{s}_K} \varphi_K(\underline{s}_K) \{1 - \varphi_K(\underline{s}_K)\} m_{c,K}(\underline{s}_K) m_{d,K}(\underline{s}_K) \\ & - \sum_{\underline{s}_K, \underline{s}'_K} I(\underline{s}_K \neq \underline{s}'_K) \varphi_K(\underline{s}_K) \varphi_K(\underline{s}'_K) \{m_{c,K}(\underline{s}_K) m_{c,K}(\underline{s}'_K) + m_{d,K}(\underline{s}_K) m_{d,K}(\underline{s}'_K)\} \\ & \left. + \sum_{\underline{s}_K, \underline{s}'_K} I(\underline{s}_K \neq \underline{s}'_K) \varphi_K(\underline{s}_K) \varphi_K(\underline{s}'_K) \{m_{d,K}(\underline{s}_K) m_{c,K}(\underline{s}'_K) + m_{c,K}(\underline{s}_K) m_{d,K}(\underline{s}'_K)\} \right] \end{aligned}$$

Assume now that $p = \frac{1}{2}$. We use (4b) to express $V_{\text{BI}}(\bar{Y}_c - \bar{Y}_d) - V_n^*(\bar{Y}_c) - V_n^*(\bar{Y}_d)$ as

$$\begin{aligned} & 2/n \left[\sum_{\underline{s}_K} \varphi_K(\underline{s}_K) \{1 - \varphi_K(\underline{s}_K)\} \{m_{c,K}(\underline{s}_K)^2 + m_{d,K}(\underline{s}_K)^2\} \right. \\ & \left. - \sum_{\underline{s}_K, \underline{s}'_K} I(\underline{s}_K \neq \underline{s}'_K) \varphi_K(\underline{s}_K) \varphi_K(\underline{s}'_K) \{m_{c,K}(\underline{s}_K) m_{c,K}(\underline{s}'_K) + m_{d,K}(\underline{s}_K) m_{d,K}(\underline{s}'_K)\} \right] \end{aligned}$$

noting that subgroup quantities under BI equal those under SCI, up to sampling variation. Combining the two expressions yields

$$\begin{aligned} V_{\text{BI}}(\bar{Y}_c - \bar{Y}_d) - V^*(\bar{Y}_c - \bar{Y}_d) &= 1/n \left[\sum_{\underline{s}_K} \varphi_K(\underline{s}_K) \{1 - \varphi_K(\underline{s}_K)\} \{m_{c,K}(\underline{s}_K) + m_{d,K}(\underline{s}_K)\}^2 \right. \\ & - \sum_{\underline{s}_K, \underline{s}'_K} I(\underline{s}_K \neq \underline{s}'_K) \varphi_K(\underline{s}_K) \varphi_K(\underline{s}'_K) \{m_{c,K}(\underline{s}_K) \\ & \left. + m_{d,K}(\underline{s}_K)\} \{m_{c,K}(\underline{s}'_K) + m_{d,K}(\underline{s}'_K)\} \right] \quad (\text{A12}) \end{aligned}$$

Algebraic calculations produce the desired result from (A12) for $w_p = w_p^{-1} = 1$, where

$w_p = \sqrt{\frac{1-p}{p}}$. The derivation for arbitrary randomization probability p formally reduces to that for $p = \frac{1}{2}$ by ‘weighting’ the subgroup means under $z = c, d$ by w_p and w_p^{-1} , respectively.

Table 1
Specification of four adaptive treatment strategies for two decision points.

	S_1	A_1	S_2	A_2	Number of randomizations
Trinary states	1	SW_1/SW_2	1	SW_2 if $A_1 = SW_1$ else SW_1	1
			2	$A_1 + AUG_1/AUG_2$	2
			3	No change	1
	2	AUG_1/AUG_2	1	SW_1/SW_2	2
			2	SW_1/SW_2	2
			3	No change	1
Binary states ^{*,†}	3	No change	1	SW_1/SW_2	1
			2	AUG_1/AUG_2	1
			3	No change	0
			1	$A_1 + 3T_1/3T_2$	2
			2	$A_1 + 2T_1/2T_2$	2
	1	$2M_1/2M_2$	1	$A_1 + 2T_1/2T_2$	2
			2	$A_1 + T_1/T_2$	2
			3	No change	2
			1	$2M_1/2M_2$	2
			2	M_1/M_2	2

^{*} $2M_i$ is twice the intensity of M_i .
[†] cT_i is c times the intensity of T_i .

Table II
Comparison of sequential causal inference (SCI) and marginal mean (MM) models when all subjects are randomized at each decision point to one of the two options.

	True variance [*]	$V^*(\bar{Y}_z)^{\dagger}$	$V_{MM}(\bar{Y}_z)^{\dagger}$
<i>Transition matrix (24)</i>			
$K^{\ddagger} = 2, \beta_z = (0, 1)$	0.034	0.033 (0.0040)	0.037 (0.0045)
$K = 2, \beta_z = (1, 2)$	0.045	0.044 (0.0044)	0.071 (0.0075)
$K = 4, \beta_z = (0, 1, 2, 3)$	0.16	0.16 (0.020)	0.34 (0.062)
$K = 4, \beta_z = (1, 2, 3, 4)$	0.29	0.27 (0.028)	0.82 (0.13)
<i>Transition matrix (25)</i>			
$K^{\ddagger} = 2, \beta_z = (0, 1)$	0.034	0.035 (0.0042)	0.038 (0.0046)
$K = 2, \beta_z = (1, 2)$	0.039	0.044 (0.0047)	0.060 (0.0068)
$K = 4, \beta_z = (0, 1, 2, 3)$	0.11	0.12 (0.014)	0.11 (0.015)
$K = 4, \beta_z = (1, 2, 3, 4)$	0.19	0.19 (0.020)	0.21 (0.027)

^{*} The 'true' value, which is calculated across 2000 simulations, was empirically the same under SCI and MM.

[†] Estimator value is the mean across 2000 simulated data sets; standard errors are given in parentheses.

[‡] Number of decision points.