

RIETI Discussion Paper Series 23-E-081

SMART-EXAM: Incorporating Participants' Welfare into Sequential Multiple Assignment Randomized Trials

WANG, Xinru Duke-NUS Medical School

DELIU, Nina University of Cambridge / Sapienza University of Rome

> NARITA, Yusuke RIETI

CHAKRABORTY, Bibhas

Duke-NUS Medical School / National University of Singapore / Duke University



The Research Institute of Economy, Trade and Industry https://www.rieti.go.jp/en/

November 2023

SMART-EXAM: Incorporating Participants' Welfare into Sequential Multiple Assignment Randomized Trials

Xinru Wang Centre for Quantitative Medicine, Duke-NUS Medical School Nina Deliu MRC - Biostatistics Unit, University of Cambridge MEMOTEF Department, Sapienza University of Rome Yusuke Narita¹ Department of Economics and Cowles Foundation, Yale University and RIETI Bibhas Chakraborty Centre for Quantitative Medicine, Duke-NUS Medical School Program in Health Services and Systems Research, Duke-NUS Medical School Department of Statistics and Data Science, National University of Singapore Department of Biostatistics and Bioinformatics, Duke University

Abstract

Dynamic Treatment Regimes (DTRs) are sequences of decision rules that recommend treatments based on patients' time-varying clinical conditions. The Sequential Multiple Assignment Randomized Trial (SMART) is an experimental design that can provide high-quality evidence for constructing optimal DTRs. In a conventional SMART, participants are randomized to available treatments at multiple stages with balanced randomization probabilities. Despite its relative simplicity of implementation and desirable performance in comparing embedded DTRs, the conventional SMART faces inevitable ethical issues including assigning many participants to the empirically inferior treatment or the treatment they dislike, which might slow down the recruitment procedure and lead to higher attrition rates, ultimately leading to poor internal and external validities of the trial results. In this context, we propose a SMART under the Experiment-as-Market framework (SMART-EXAM), a novel SMART design that holds the potential to improve participants' welfare by incorporating their preferences and predicted treatment effects into the randomization procedure. We describe the steps of conducting a SMART-EXAM and evaluate its performance compared to the conventional SMART. The results indicate that the SMART-EXAM can improve the welfare of the participants enrolled in the trial, while also achieving a desirable ability to construct an optimal DTR when the experimental parameters are suitably specified. We finally illustrate the practical potential of the SMART-EXAM design using data from a SMART for children with attentiondeficit/hyperactivity disorder (ADHD).

Keywords: Sequential Multiple Assignment Randomized Trials; Dynamic Treatment Regimes; Experiment-as-Market; Preference-adaptive design; Response-adaptive design

The RIETI Discussion Paper Series aims at widely disseminating research results in the form of professional papers, with the goal of stimulating lively discussion. The views expressed in the papers are solely those of the author(s), and neither represent those of the organization(s) to which the author(s) belong(s) nor the Research Institute of Economy, Trade and Industry.

¹This study is conducted as a part of research studies at the Research Institute of Economy, Trade and Industry (RIETI). The draft of this paper was presented at a DP seminar of the Research Institute of Economy, Trade and Industry (RIETI). I would like to thank participants of the RIETI DP Seminar for their helpful comments.

1 Introduction

The management of chronic and relapsing diseases requires adjusting treatments at different time points based on previous treatment histories and disease status to improve the final outcomes of interest. Dynamic treatment regimes (DTRs) are sequences of decision rules that can formalize the patient-centered disease management model designed to guide clinicians in prescribing the right treatments to the right patients. The advantages of DTRs over fixed treatments have been recognized by researchers in various health domains, from the perspectives of maximizing treatment effectiveness, minimizing side effects, and improving cost-effectiveness. The sequential multiple assignment randomized trial (SMART), an experimental design with multiple randomization stages, is considered the gold-standard design for constructing optimal DTRs (Murphy, 2005; Wang and Chakraborty, 2023). Figure 1 illustrates a SMART for children with attention-deficit/hyperactivity disorder (ADHD) (Pelham Jr et al., 2016). At stage 1, children with ADHD were randomized to either low-intensity behavioral modification (BMOD) or low-dose oral methamphetamine (MEDS). After eight weeks of treatment, the response status of each child was assessed using the Impairment Rating Scale and an individualized list of target behaviors on a monthly basis. If the children showed insufficient response to the treatment, they were re-randomized to either intensify the initial treatment or augment with another treatment. Responders continued with the initial treatment and got re-randomized if the disease conditions deteriorated at any time after eight weeks. In this SMART, there were four embedded DTRs, one of which is: "begin with low-intensity BMOD and add MEDS at stage 2 for non-responders."

[Figure 1 about here.]

Numerous methods have been established to analyze data from SMARTs, with the aim of estimating stage-specific treatment effects, comparing the embedded DTRs, or constructing more tailored DTRs (Oetting et al., 2011; Nahum-Shani et al., 2012). From the design perspective, the conventional SMART (henceforth referred to as the SMART where there is no confusion), i.e., SMART with balanced randomization probabilities across treatments at each stage, can maximize the ability to compare the embedded DTRs (Murphy, 2005). However, with various data sources available, e.g., electronic health records, cohort studies, patient databases, and previous trials, for estimating treatment effects, it has become debatable whether it is ethical to evenly randomize participants without considering the potential individualized treatment effects. Furthermore, in recent years, a broader trend towards patient-centered care has emerged, with a growing emphasis on patient preferences and shared decision making (Fukami, 2023). Under this framework, healthcare providers and patients exchange information about patients' disease-related characteristics and the benefits and side effects of each treatment, and finally achieve a consensus about which treatment(s) should be prescribed. This paradigm shift is grounded in the principle of patient autonomy, with the aim of promoting patients' participation in healthcare decision-making. In view of this paradigm shift, it may be unreasonable to evenly randomize participants ignoring their preferences in SMARTs altogether. Particularly noteworthy is the fact that the treatment components in a SMART are often available in routine care which may make participants hesitate to join the trial due to concerns about lower probabilities of receiving their preferred treatment by entering into the trial (Edwards and Braunholtz, 2000). Ignoring the above issues may exert a negative impact on the recruitment procedure and demotivate participants from sticking to their assigned treatments, which may ultimately reduce the power and internal and external validities of the trial results (Gaines and Kuklinski, 2011).

The above ethical issues are also present in randomized controlled trials (RCTs). The response-adaptive randomization (RAR) (Eisele, 1994; Rosenberger and Hu, 2004; Robertson et al., 2023) was proposed to allow for potential adjustments of randomization probabilities, in accordance with accumulating information about the treatment performance. In addition,

a lot of work has been done to either incorporate participants' preferences during randomization or estimate average choice-specific treatment effects (ACTE), i.e., the conditional treatment effect among those who would choose a particular treatment option. Specifically, two-stage randomized designs (Rücker, 1989) first randomize participants into either the randomization group or the choice group. Participants in the randomization group will be randomized again to the available treatments, while participants in the choice group will receive their preferred treatment. By randomizing twice, such designs enable researchers to estimate ACTE (Knox et al., 2019) and increase the probability of allocating participants to their preferred treatment as compared to conventional RCTs. In addition, fully randomized preference designs (Torgerson et al., 1996) first elicit participants' preferences before randomization, and then use the elicited preferences to estimate ACTE. One of their limitations is that there may be discrepancies between participants' stated preferences and actual choices, and it may be unethical to continue with balanced randomization probabilities after asking about their preferences. In partially randomized preference designs, participants with strong preferences are given their preferred treatment, while those with mild preferences are randomized to available treatments. While such designs have the potential to recruit more participants with strong preferences who would have refused to participate had they got randomized, the estimated average treatment effects (ATEs) are restricted among those with mild preferences instead of a wider population (Walter et al., 2017).

There have been several SMARTs that incorporated participants' preferences during the randomization procedure (Fava et al., 2003; McKay et al., 2015). The Sequenced Treatment Alternatives to Relieve Depression (STAR*D) was a multisite, multi-stage, first-generation SMART for patients with major depressive disorder (Fava et al., 2003). In STAR*D, participants were randomized within their preferred treatment categories (switch or augment) during the second and third stages. From the analysis perspective, there have been some

works (Wang et al., 2012; Johnson et al., 2021; Xu et al., 2022) to construct composite outcomes accounting for both treatment efficacy and other factors, e.g., toxicity and quality of life, to evaluate DTRs. The parameters in the outcome regarding these factors may be chosen based on patients' preferences, e.g., whether patients prefer higher treatment effects, lower side effects, or lower costs. All of these works demonstrate the paramount importance of incorporating preferences in SMARTs.

Furthermore, attempts have been made to incorporate predicted treatment effects in SMARTs. Cheung et al. (2015) proposed a SMART with adaptive randomization (SMART-AR) using Q-learning to determine the randomization probabilities in favor of superior treatments based on the complete data trajectories from previous participants in the trial. Wang et al. (2022) introduced a response-adaptive SMART (RA-SMART) that employs a framework akin to the "play the winner" rule, i.e., the inferior treatment at stage 1 will have a lower randomization probability at stage 2. Wu et al. (2023) presented a SMART with interim monitoring, where early termination is permitted if there is sufficient evidence of treatment efficacy. Nevertheless, none of the aforementioned designs simultaneously incorporated participants' preferences and treatment effects into the randomization procedure.

Recently, Narita (2021) proposed an extension of the RCT — the Experiment-as-Market (EXAM) design, which allows for incorporating both participants' preferences and treatment effects into the randomization procedure, while maintaining robust inferences comparable to RCTs. Inspired by this, the current paper proposes a novel SMART design under the Experiment-as-Market framework (SMART-EXAM), which has the potential to improve participants' welfare in terms of both the participants' preferences and their final outcomes. By taking treatment effects and participants' preferences into account to individualize the randomization probabilities, the SMART-EXAM mimics a shared decision-making process

between patients and healthcare providers in a clinical trial setting, which enables participants to actively participate in the randomization procedure.

The remainder of the article is divided into five sections. Section 2 illustrates the structure of a generic two-stage SMART used throughout this paper and presents analysis methods to compare embedded DTRs in a SMART. Section 3 presents the detailed procedure of conducting a SMART-EXAM and its theoretical properties. In Section 4, we conduct a simulation study to evaluate the empirical performance of the SMART-EXAM compared to the SMART. In Section 5, we demonstrate the practical potential of the SMART-EXAM using data from a SMART ADHD study. The last section comes with conclusions and discussions about the SMART-EXAM.

2 Set up and Notation

To facilitate the exposition, we focus on a two-stage SMART that is consistent with the SMART ADHD study (Pelham Jr et al., 2016). The observed trajectory for the *i*-th participant is denoted by $(O_{1i}, A_{1i}, O_{2i}, A_{2i}, Y_i)$, which is assumed to be independent and identically distributed. O_{ti} (t = 1, 2) denotes the vector of covariates obtained prior to treatment at stage *t*. Among covariates O_{2i} , R_i is the indicator for intermediate response status, with $R_i = 1/0$ for responders and non-responders, respectively. A_{ti} is the treatment indicator at stage *t*, with $A_{1i} \in \{-1, 1\}$, $A_{2i} \in \{-1, 1\}$ for non-responders, and $A_{2i} = 0$ for responders. Y_i is the final continuous outcome and without loss of generality, we assume that higher values of Y_i are preferred. Define $H_{ti} \in \mathcal{H}_t$ as the history data of the *i*-th participant at stage *t*, with $H_{1i} = O_{1i}$, $H_{2i} = (O_{1i}, A_{1i}, O_{2i})$, where \mathcal{H}_t is the space of possible histories at stage *t*. Under the SMART shown in Figure 1, the *j*-th two-stage DTR in the space of DTRs of interest \mathcal{D} is denoted as $d_j = (d_{j1}, d_{j2}), j = 1, \ldots, J$, where $d_{j1}, d_{j2} \in \{-1, 1\}$, and J is the total number of embedded DTRs. For example, $d_1 = (1, 1)$ means that, first treat participants with $A_1 = 1$, if they do not respond, switch to $A_2 = 1$, otherwise continue with

the initial treatment. Under the Neyman-Rubin causal inference framework (Rubin, 1974), the potential outcome under DTR $d_j = (d_{j1}, d_{j2})$ is denoted as

$$Y^{d_j} = R^{d_{j1}} Y^{s(d_{j1},0)} + (1 - R^{d_{j1}}) Y^{s(d_{j1},d_{j2})},\tag{1}$$

where $s(d_{j1}, 0)$ denotes the treatment sequence of receiving $A_1 = d_{j1}$, responding, and continuing with the initial treatment, while $s(d_{j1}, d_{j2})$ denotes the treatment sequence of receiving $A_1 = d_{j1}$, not responding, and switching to $A_2 = d_{j2}$. Here we use the prefix s to differentiate the notations for DTR d_j and the treatment sequence. $Y^{s(d_{j1},0)}$ and $Y^{s(d_{j1},d_{j2})}$ are the potential outcomes for participants with treatment sequences $s(d_{j1}, 0)$ and $s(d_{j1}, d_{j2})$, respectively. The expected outcome of DTR d_j is $\mu_{d_j} = E(Y^{d_j}) = \pi_{d_{j1}}\mu_{s(d_{j1},0)} + (1 - \pi_{d_{j1}})\mu_{s(d_{j1},d_{j2})}$, where $\pi_{d_{j1}}$ is the response rate for treatment d_{j1} at stage 1; $\mu_{s(d_{j1},0)}$ and $\mu_{s(d_{j1},d_{j2})}$ are the expected outcomes of those with treatment sequences $s(d_{j1}, 0)$ and $s(d_{j1}, d_{j2})$, respectively. The optimal DTR is defined as the DTR with the highest expected outcome μ_{d_j} , i.e., $d^* = \arg \max \mu_{d_j}$.

In this paper, the primary goal is to select an optimal DTR d^* among the embedded DTRs in a SMART. Although there are various methods for estimating DTR means, e.g., Gcomputation and augmented inverse probability weighting (AIPW), we mainly focus on the inverse probability weighting (IPW) due to both its unbiasedness and ease of implementation in the context of the SMART-EXAM. Under usual causal assumptions in the Neyman-Rubin causal framework, the IPW estimator for the expected outcome of DTR $d_j = (d_{j1}, d_{j2})$ is defined as

$$\hat{\mu}_{d_j} = \frac{\sum_{i=1}^{N} W_i^{d_j} Y_i}{\sum_{i=1}^{N} W_i^{d_j}},\tag{2}$$

where N is the total sample size, $W_i^{d_j} = \frac{I(A_{1i}=d_{j1})}{p_{1,d_{j1},i}} \left\{ R_i + \frac{(1-R_i)I(A_{2i}=d_{j2})}{p_{2,d_{j2},i}} \right\}$ is the weight of the *i*-th individual for DTR $d_j = (d_{j1}, d_{j2}), p_{1,d_{j1},i} = \Pr(A_{1i} = d_{j1}), \text{ and } p_{2,d_{j2},i} = \Pr(A_{2i} = d_{j2}|R_i = 0, A_{1i} = d_{j1})$. The variance estimator of the DTR mean estimator (Oetting et al.,

2011) is

$$\widehat{\operatorname{Var}}(\hat{\mu}_{d_j}) = \frac{\sum_{i=1}^{N} (W_i^{d_j} (Y_i - \hat{\mu}_{d_j}))^2}{N^2}.$$
(3)

The estimated optimal DTR $\hat{\boldsymbol{d}}^*$ is thus defined as $\hat{\boldsymbol{d}}^* = \underset{\boldsymbol{d}_j}{\operatorname{arg\,max}} \hat{\mu}_{\boldsymbol{d}_j}$.

3 SMART with the Experiment-as-Market Framework (SMART-EXAM)

As an extension of RCTs, the EXAM aims to strike a balance between making accurate inferences for future patients and improving the welfare of enrolled participants through an imaginary centralized market (Narita, 2021). One can view the EXAM as a shared decisionmaking randomization procedure where healthcare providers rely on existing evidence about treatment effects to determine the most effective treatment, while participants indicate their preferences based on previous treatment experience and information provided by professionals about potential benefits and side effects of each treatment. The EXAM can generate randomization probabilities that take into account the perspectives of both healthcare providers and participants and achieve a so-called consensus about the randomization probabilities. It has been proved that the EXAM can improve participants' welfare while ensuring valid inferences about the ATEs, on the grounds that the EXAM can be seen as a special case of stratified randomized trials based on predicted treatment effects and preferences.

Building upon the work of Narita (2021), the current paper proposes a novel SMART design - the SMART-EXAM - to simultaneously incorporate participants' preferences and treatment benefits into the randomization procedure. Note that the embedded DTRs in a SMART are often stepped-up treatment strategies, i.e., for non-responders, stepped-up treatments are provided to "rescue" the initial treatment. In this case, it is critical to improve the allocation at stage 2. Furthermore, participants tend to have preferences for stage-2 treatments, as is shown in a pilot study for adolescent depression (Gunlicks-Stoessel et al., 2016). Against this

backdrop, the current paper focuses on SMART-EXAMs with individualized randomization at stage 2, while maintaining balanced randomization at stage 1.

3.1 Key Definitions

To facilitate discussion, we first provide basic definitions for a SMART-EXAM. Among nonresponders to the initial treatment a_1 , we define:

DEFINITION 1 (*Preferences*): Let $\Lambda_i \in \{0, 1\}$ denote the stage-2 preference indicator of the *i*-th participant, with $\Lambda_i = 1$ and $\Lambda_i = 0$ corresponding to the cases where the participant prefers $A_2 = 1$ and $A_2 = -1$, respectively.

DEFINITION 2 (Treatment capacity and Treatment demand): Let $C_{a_2|a_1}$ be the capacity for a_2 among non-responders to a_1 , such that $\sum_{a_2 \in \{-1,1\}} C_{a_2|a_1} = N_{a_1}^{nr}$, where $N_{a_1}^{nr}$ is the number of non-responders to a_1 . The treatment demand $D_{a_2|a_1}$ is defined as the total number of non-responders to a_1 who prefer a_2 , i.e., $D_{a_2|a_1} = \sum_{i=1}^{N_{a_1}^{nr}} \{I(a_2 = 1)\Lambda_i + I(a_2 = -1)(1 - \Lambda_i)\}$. Define the excess-demand and the oversupplied stage-2 treatment as $a_2^e = \{a_2 \in \{-1,1\} : D_{a_2|a_1} \ge C_{a_2|a_1}\}$ and $a_2^o = \{a_2 \in \{-1,1\} : D_{a_2|a_1} < C_{a_2|a_1}\}$, respectively.

The choice of treatment capacities depends on various factors (Dumville et al., 2006). For example, allocating less capacity to the more expensive treatment can reduce trial costs, while allocating higher capacity to the treatment with a higher drop-out rate and variance can enhance the overall statistical efficiency. In addition, considerations may also involve assigning more capacity to the treatment on which researchers may be interested in gaining more information, e.g., side effects or treatment experience and assigning more capacity to the treatment that most participants prefer.

DEFINITION 3 (Individualized treatment effects): Under the Neyman-Rubin causal inference framework, the effect of $A_2 = a_2^e$ on the outcome compared with $A_2 = a_2^o$ conditional on the history \boldsymbol{H}_{2i} for participants $i = 1, \ldots, N_{a_1}^{nr}$ is denoted as $\zeta_{2,a_2^e,i} = E[Y_i|\boldsymbol{H}_{2i}, A_{2i} =$ a_2^e] – $E[Y_i|\boldsymbol{H}_{2i}, A_{2i} = a_2^o]$. For the simplicity of notations, we use ζ_i for such treatment effects where there is no confusion. For example, the outcome model can be specified as

$$Q_2(\boldsymbol{H}_{2i}, A_{2i}; \boldsymbol{\gamma}_2, \boldsymbol{\alpha}_2) = E[Y_i | \boldsymbol{H}_{2i}, A_{2i}] = \boldsymbol{\gamma}_2^T \boldsymbol{H}_{20i} + \boldsymbol{\alpha}_2^T \boldsymbol{H}_{21i} A_{2i}, \tag{4}$$

where \boldsymbol{H}_{20i} and \boldsymbol{H}_{21i} are potentially different features of \boldsymbol{H}_{2i} , and \boldsymbol{H}_{21i} is the vector of covariates in \boldsymbol{H}_{2i} that are believed to interact with the treatment. The individualized treatment effect $\hat{\zeta}_i$ is $2\hat{\boldsymbol{\alpha}}_2^T \boldsymbol{H}_{21i}$ if $a_2^e = 1$, or $-2\hat{\boldsymbol{\alpha}}_2^T \boldsymbol{H}_{21i}$ if $a_2^e = -1$, where $\hat{\boldsymbol{\alpha}}_2^T$ is the estimated coefficients based on previous data. Note that when estimating treatment effects using observational data, careful attention should be paid to adjust for potential confounders. In addition to Q-learning with linear regression, other advanced statistical methods, e.g., random forest, regression trees, and boosting, etc. (Knaus et al., 2021), can be utilized to estimate ζ_i . Here we use the model in Eq. (4) only for the purpose of illustration.

Treatment effects and preferences are two complementary measures of participants' welfare, which are sometimes found to be correlated through some common predictors, e.g., age and previous treatment experience. For instance, in the SMART ADHD study, Nahum-Shani et al. (2012) found that non-responders with lower adherence have higher effects of augmentation over intensification compared to those with higher adherence. It is also reasonable to assume that non-responders with lower adherence are more likely to prefer augmentation than those with higher adherence. As a result, there is a positive association between treatment effects and preferences. It aligns with the shared decision-making procedure with less tradeoff between preferences and treatment effects as perceived by healthcare providers.

DEFINITION 4 (Utility): Define the utility of the *i*-th participant as $u_i = p_{2,1,i} \Lambda_i + (1 - p_{2,1,i})(1 - \Lambda_i)$.

To appreciate the implications of the utility, we introduce the indicator for whether the *i*-th participant receives the preferred treatment or not, denoted by $K_i = I(A_{2i} = 1)\Lambda_i + I(A_{2i} = -1)(1 - \Lambda_i)$. The probability of receiving the preferred treatment, i.e., the expected value of

 K_i , conditional on participants' indicated preferences, is thus denoted as $u_i = E[K_i|\Lambda_i] = p_{2,1,i} \Lambda_i + (1 - p_{2,1,i})(1 - \Lambda_i).$

Utility holds significant importance for participants when deciding whether or not to participate in the trial. As stated in Section 1, participants may refuse to participate due to concerns about a relatively lower probability of receiving their preferred treatment by entering into the trial. However, some "altruistic" participants may accept making small sacrifices and participating to contribute to future patients. The higher the probability of receiving their preferred treatment, the less is altruism required and more participants may agree to participate in SMARTs (Edwards and Braunholtz, 2000).

3.2 Procedure of conducting a SMART-EXAM

Figure 2 graphically illustrates a SMART-EXAM that individualizes the stage-2 randomization probabilities, and the corresponding detailed algorithm is available in Web Appendix A. The stage-1 randomization is the same as that in the SMART, while at the intermediate decision point, the preference indicators Λ_i are collected and treatment effects $\hat{\zeta}_i$ are estimated based on previous trials or observational studies. In the following procedure, the $\hat{\zeta}_i$ are normalized so that they have a mean of 0 and a standard deviation of 1. To produce less variable randomization probabilities, the normalized treatment effects $\hat{\zeta}_i$ should be discretized by one of the binning strategies, e.g., creating contiguous intervals with equal frequencies and specifying the categorized value of treatment effect $\tilde{\zeta}_i$ as the mean value of the interval.

[Figure 2 about here.]

Subsequent to estimating treatment effects, each non-responder is given a budget m = 1. The same budget for each participant is to ensure giving equal importance to every participant's welfare. Define Ψ_i as the treatment price of a unit randomization probability of a_2^e for participant *i*; there exist $\eta \in (-1, 0)$ and $\beta \in \mathcal{R}$ such that $\Psi_i = \eta \tilde{\zeta}_i + \beta$. Note

that β and η are common parameters for the price function; the individualization of price Ψ_i is induced by the heterogeneous treatment effect $\tilde{\zeta}_i$. Intuitively, $\tilde{\zeta}_i < \tilde{\zeta}_{i'}$ corresponds to the cases where the *i'*-th participant will benefit more when $\tilde{\zeta}_{i'} > 0$ (or lose less when $\tilde{\zeta}_{i'} < 0$) from treatment a_2^e than the *i*-th participant. Thus it is more reasonable to give the *i'*-th participant a higher randomization probability for a_2^e than the *i*-th participant when the treatment resources are limited. This can be accomplished by a negative value of η , which ensures that, $\Psi_i > \Psi_{i'}$ if $\tilde{\zeta}_i < \tilde{\zeta}_{i'}$, i.e., the price of a unit randomization probability for a_2^e is lower for those with higher treatment effects, such that they can "buy" more randomization probabilities for a_2^e compared to those with lower treatment effects. The values of β depend on treatment capacities, which will be described later.

The randomization probabilities are derived by solving the utility maximization function subject to budget constraints. For participant i:

$$\hat{p}_{2,a_2^e,i} = \underset{p_{2,a_2^e,i} \in \mathcal{P}}{\operatorname{arg\,max}} u_i \quad s.t. \ p_{2,a_2^e,i} \ \Psi_i \leqslant m,$$
(5)

where $\mathcal{P} \equiv \{p_{2,a_{2}^{e},i} | p_{2,a_{2}^{e},i} \in [0,1]\}$, and $p_{2,a_{2}^{e},i} \Psi_{i}$ is the expected expense of "buying" randomization probabilities for treatment a_{2}^{e} . The underlying rationale of Eq. (5) is to maximize the utility, i.e., the probability of receiving the preferred treatment, subject to the budget constraint, with the price determined by the predicted treatment effects. The budget constraint takes predicted treatment effects into the optimization problem and ensures that the limited treatment resources favor those who can benefit more from the treatment.

To ensure that a non-null subset of participants follow each embedded DTR, i.e., the positivity assumption holds for valid inferences, $\epsilon \in [0, \bar{\epsilon}]$ is introduced to keep the randomization probabilities within the range $[\epsilon, 1 - \epsilon]$, where $\bar{\epsilon} = \min_{a_2} p_{2,a_2}^0$, and $p_{2,a_2}^0 = \frac{C_{a_2|a_1}}{N_{a_1}^{a_1}}$ is the non-individualized randomization probability for a_2 with treatment capacity $C_{a_2|a_1}$. The final stage-2 randomization probability is $\tilde{p}_{2,a_2,i} = (1 - q)\hat{p}_{2,a_2,i} + qp_{2,a_2}^0$, where $q \equiv \inf_{i,a_2} \{q' \in [0,1]|(1 - q')\hat{p}_{2,a_2,i} + q'p_{2,a_2}^0 \in [\epsilon, 1 - \epsilon]\}$. Alternatively, other methods such as trimming the extreme probabilities to a predetermined value can also avoid violating the positivity assumption. The values of β in the price function are adjusted through an iterative procedure shown in Figure 2, to ensure that the randomization probabilities derived from the above procedure meet the capacity constraints: $\sum_{i=1}^{N_{a_1}^{nr}} \tilde{p}_{2,a_2,i} = C_{a_2|a_1}$ for each $a_2 \in \{-1, 1\}$, i.e., the expected number of non-responders allocated to a_2 under the current randomization mechanism equals to the capacity of a_2 ; thus, the design satisfies a system-level capacity constraint. Note that in Figure 2, interval l can be interpreted as the step size for updating the value of β . $(d_{a_2}^2 + d_{a_2}^2)/(N_{a_1}^{nr})$ is compared to κ to assess whether the sum of the square differences between the expected number of participants and the capacity for each treatment, i.e., $d_{a_2}^2 + d_{a_2}^2$, is within a desired range $[0, \kappa N_{a_1}^{nr}]$ or not.

By the above procedure, participants with the same Λ_i and $\tilde{\zeta}_i$ will be given the same randomization probabilities as they have the same utility and budget constraint. In other words, non-responders can be divided into B groups indexed by $G_i \in \{1, \ldots, B\}$ based on Λ_i and $\tilde{\zeta}_i$, such that for those in the group $g = 1, \ldots, B$, $\tilde{p}_{2,a_2,i}$ is the same and is denoted by $p_{2,a_2|a_1,g}$. The DTR means can be estimated by Eq. (2). The procedure for individualizing stage-1 randomization probabilities in a SMART-EXAM is available in Web Appendix B.

3.3 Theoretical properties of the SMART-EXAM

The following are three theoretical properties of the SMART-EXAM, and the corresponding proofs are provided in Web Appendix C.

PROPERTY 1: When the participants' preferences and predicted treatment effects are not of concern, a SMART-EXAM with capacity $C_{a_2|a_1}$ can be reduced to a SMART with nonindividualized randomization probabilities $p_{2,a_2}^0 = \frac{C_{a_2|a_1}}{N_{a_1}^{nr}}$, by setting $\Lambda_i = \Lambda_j$ and $\tilde{\zeta}_i = \tilde{\zeta}_j$ for all i and j $(j \neq i)$ among non-responders to the initial treatment a_1 .

PROPERTY 2: Under the three causal assumptions in the Neyman-Rubin causal frame-

work: 1) sequential exchangeability assumption (SEA); 2) consistency assumption (CA); and 3) positivity assumption (PA), which are detailed in Web Appendix C, the IPW estimator for the value of DTR \mathbf{d}_j in a SMART-EXAM is a consistent estimator of the expected outcome of \mathbf{d}_j , i.e., $\hat{\mu}_{\mathbf{d}_j} \xrightarrow{p} \mu_{\mathbf{d}_j}$.

PROPERTY 3: The large-sample distribution of $\hat{\mu}_{d_j}$ for DTR d_j is $\sqrt{N}(\hat{\mu}_{d_j} - \mu_{d_j}) \xrightarrow{d} N(0, \sigma_{d_j}^2)$, where

$$\sigma_{\boldsymbol{d}_{j}}^{2} = \frac{\pi_{1,d_{j1}}}{p_{d_{j1}}} \{ \sigma_{s(d_{j1},0)}^{2} + (\mu_{\boldsymbol{d}_{j}} - \mu_{s(d_{j1},0)})^{2} \} + \sum_{g} \left[\frac{(1 - \pi_{d_{j1}})\pi_{g|d_{j1}}}{p_{1,d_{j1}}p_{2,d_{j2}|a_{1},g}} \times \{ \sigma_{s(d_{j1},g,d_{j2})}^{2} + (\mu_{\boldsymbol{d}_{j}} - \mu_{s(d_{j1},g,d_{j2})})^{2} \} \right],$$

$$(6)$$

with $\mu_{s(d_{j1},g,d_{j2})} = E[Y|A_1 = d_{j1}, R = 0, G = g, A_2 = d_{j2}], \ \sigma^2_{s(d_{j1},g,d_{j2})} = \operatorname{Var}(Y|A_1 = d_{j1}, R = 0)$ $0, G = g, A_2 = d_{j2}), \ \mu_{s(d_{j1},0)} = E[Y|A_1 = d_{j1}, R = 1], \ \sigma^2_{s(d_{j1},0)} = \operatorname{Var}(Y|A_1 = d_{j1}, R = 1),$ $and \ \pi_{g|d_{j1}} = \Pr(G = g|A_1 = d_{j1}, R = 0)$

4 Simulations

Using 500 simulation replicates, we compare the performance of the SMART-EXAM with the SMART in different settings detailed in Section 4.1.

4.1 Data generation

Table 1 provides the parameter specification for generating data from SMART designs. As discussed in Section 3.1 that there may be positive or negative associations between treatment effects and preferences, we consider two simplified scenarios for generating participants' preferences. The first assumes a negative association between preferences Λ_i and treatment effects ζ_i , i.e., participants with higher effects of $A_2 = a_2^e$ are less likely to prefer $A_2 = a_2^e$ compared to those with lower effects. The second assumes a positive association between preference Λ_i and treatment effect ζ_i . It is important to note that our approach to generating preference indicators does not imply that the SMART-EXAM requires a specific association between these two measures. Rather, it serves as a means to explore the performance of the SMART-EXAM with different degrees of trade-off between preferences and treatment effects. We reemphasize that, since our goal was not to capture the full complexity of participants' preferences, we only assume a simplified association between these two measures in simulation studies, which is sufficient to provide valuable insights.

Furthermore, we consider different parameter sets for ϵ and η in SMART-EXAMs detailed in Table 1. Recall that ϵ controls the range of the randomization probabilities and η is the coefficient in the price function. Parameter sets 1) - 3) are to explore the impact of ϵ on the performance of SMART-EXAMs, and SMART-EXAMs with smaller values of ϵ are expected to generate more extreme probabilities for improving participants' welfare; while parameter sets 2) and 4) - 6) are to investigate the impact of η on the performance of SMART-EXAMs. The predicted treatment effects $\hat{\zeta}_i$ are derived based on pilot SMARTs with different sample sizes $N^p = 100, 200, 300$ simulated in the same manner as above. These pilot SMARTs are to investigate the sensitivity of the SMART-EXAM to the quality of the external information. We also consider different sample sizes for the full-scale SMART N = 200, 300, 400, 500 and treatment capacities $C_{1|a_1}$, including $0.5N_{a_1}^{nr}$, $0.6N_{a_1}^{nr}$, and $0.7N_{a_1}^{nr}$, denoted by C = 0.5, C = 0.6, and C = 0.7 respectively for the simplicity of illustration. Furthermore, we include a SMART-EXAM with adaptive randomization named SMART-AR-EXAM and assume equal treatment capacity, i.e., C = 0.5. In a SMART-AR-EXAM, the first $N_{\rm min} = 100$ participants will get randomized using balanced randomization. The data from these participants instead of a pilot SMART are then used to estimate treatment effects and to generate the randomization probabilities for the remaining participants. The SMART-AR-EXAM and the SMART-EXAM resemble the duality between adaptive and fixed RCTs, where the former is an adaptive design that allows for updating randomization probabilities based on the accumulating data, while the latter generates fixed but unbalanced randomization probabilities based on external information.

[Table 1 about here.]

4.2 Metrics for the design performance

The true value of μ_{d_j} is approximated via IPW based on a simulated SMART of size 100,000. The DTR mean $\hat{\mu}_{d_j}$ and its variance $\hat{\sigma}_{d_j}^2$ are estimated by Eq. (2) and Eq. (3), respectively. The metrics to compare these designs can be divided into two categories: the first is the welfare of enrolled participants, including the participants' average outcome $\bar{Y} = \frac{1}{N^{nr}} \sum_{i=1}^{N^{nr}} Y_i$, where N^{nr} is the number of non-responders, the average probability of receiving the preferred treatment $\bar{u} = \frac{1}{N^{nr}} \sum_{i=1}^{N^{nr}} u_i$, and the number of participants in DTR d_j denoted by N_{d_j} . Note that we only focus on non-responders when comparing participants' welfare because responders only have one treatment option at stage 2. The second is the learning ability of SMART designs, evaluated by whether they correctly estimate the true optimal DTR or not denoted by $I(\hat{d}^* = d^*)$, and the true value of the estimated optimal DTR $\mu_{\hat{d}^*}$, which can be interpreted as the expected outcome if the estimated optimal DTR \hat{d}^* learned from the current SMART is given to the entire population. These metrics are collected in each simulation replicate and averaged over to get the expected value.

4.3 Simulation results

This subsection gives the simulation results for the setting with a pilot SMART of size $N^p = 200$. Results for other settings with pilot SMARTs of size $N^p = 100,300$ are provided in Web Appendix D. Figure 6 shows that, when setting equal treatment capacity for each treatment, the SMART has better learning ability represented by $E[\mu_{\hat{d}^*}]$, i.e., the Monte Carlo estimate of the true value of the estimated optimal DTR, compared to the SMART-EXAM C = 0.5, except for the cases with sample size N = 500 and a positive association between treatment effects and preferences, where SMART-EXAM with C = 0.5 performs slightly better than the SMART. However, as ϵ increases, the learning abilities of SMART-EXAM with C = 0.5 become closer to the SMART. A possible explanation might be that a lower value of ϵ corresponds to more extreme randomization probabilities, resulting in higher variances of estimated DTR means and thus less accurate estimation of optimal DTRs. As the sample size N increases, the learning abilities for all these designs improve.

In addition, all the SMART-EXAM designs demonstrate a substantial improvement in the average probability of receiving the preferred treatment $E[\bar{u}]$ and final outcome $E[\bar{Y}]$ compared to the SMART, which is as expected given that the SMART does not take into account the participants' preferences and treatment effects during randomization. While the SMART-AR-EXAM performs better in selecting the true optimal DTR compared to the SMART-EXAM with C = 0.5 in most scenarios, it offers fewer improvements in terms of $E[\bar{u}]$ and $E[\bar{Y}]$. This is reasonable given that some participants in a SMART-AR-EXAM are randomized using balanced randomization probabilities instead of individualized randomization based on treatment effects and preferences.

[Figure 3 about here.]

One can also see from Figure 6 that when C increases from 0.5 to 0.7, $E[\mu_{\hat{d}^*}]$ and $E[\bar{Y}]$ increase, as more participants are randomized to the true superior treatment $A_2 = 1$. However, in reality, the true superior treatment is unknown and the treatment capacity C should be specified based on various factors mentioned in Section 3.1. The coefficient η has a relatively smaller impact on the performance of SMART-EXAMs. Nonetheless, as the absolute value of η increases, there is a slight upward trend in the mean outcomes $E[\bar{Y}]$. This finding aligns with the underlying concept of η , where higher values of η amplify the predicted treatment effect's impact on the randomization probabilities, resulting in better performance in improving the outcome. Moreover, SMART-EXAMs demonstrate higher mean outcomes $E[\bar{Y}]$ in scenarios with a positive association between treatment effects and preferences compared to those with a negative association, which is not unanticipated given that the tradeoff between preferences and treatment effects of a positive association is expected to be less than that of a negative association. We also provide results under settings with other pilot SMARTs in Web Figures 1-2 in Web Appendix D, which show similar patterns to the aforementioned findings.

In addition to IPW estimators, we explore using AIPW for estimating DTR means. Furthermore, we employ other metrics to evaluate the operating characteristics of the SMART-EXAM designs. More specifically, we consider 1) the probabilities of being selected as the optimal DTR for each DTR d_j by IPW and AIPW, denoted as $\Pr(d_j = \hat{d}^{*IPW})$ and $\Pr(d_j = \hat{d}^{*AIPW})$; 2) the average estimated value of DTR d_j by IPW and AIPW, denoted as $E[\hat{\mu}_{d_j}^{IPW}]$ and $E[\hat{\mu}_{d_j}^{AIPW}]$, and 3) the average number of participants assigned to DTR d_j , i.e., \overline{N}_{d_j} for the scenario with sample size N = 200, $N^p = 200$, $\eta = -1$ and $\epsilon = 0.1, 0.2, 0.3$. Web Tables 1-3 in Web Appendix D show that, in all these SMART designs, $E[\hat{\mu}_{d_j}^{IPW}]$ and $E[\hat{\mu}_{d_j}^{AIPW}]$ are close to the true value of the corresponding DTR; thus, all of these designs achieve the desired performance of estimating DTR means. Furthermore, when $\epsilon = 0.3$, all SMART-EXAMs exhibit better performance than the SMART in terms of selecting the true optimal DTR, except that the SMART-EXAM with C = 0.5 performs slightly worse than the SMART when there is a negative association between treatment effects and preferences; the probability of selecting the true optimal DTR for the SMART-EXAM with C = 0.5 is 0.812, whereas for the SMART it is 0.824. The SMART-AR-EXAM demonstrates a comparable performance to the SMART in terms of selecting optimal DTRs and estimating DTR means, thus it may serve as an alternative design that can incorporate participants' welfare when there is no prior information about treatment effects.

5 Empirical Application

In this section, we illustrate the merits of the SMART-EXAM using the aforementioned SMART ADHD study as a previous/pilot SMART. Investigators from the University of Michigan provide manually generated data based on this two-stage ADHD SMART with a sample size N = 150, comprised of 51 responders and 99 non-responders (Almirall et al., 2011). The baseline covariates include O_{11} : the indicator for oppositional defiant disorder (ODD) diagnosis, coded as 1/0, O_{12} : continuous ADHD score reflecting ADHD symptoms at the end of the previous school year, O_{13} : the indicator for medication prior to first-stage treatments, coded as 1/0, and O_{14} : the indicator for whether the race is White or not, coded as 1/0. The intermediate potential tailoring variables include O_{21} : the number of months until non-response and O_{22} : the indicator for adherence to the stage-1 treatments, coded as 1/0 for higher and lower adherence. Imagine that researchers wish to further explore the effectiveness of the embedded DTRs for children with ADHD, and plan to conduct a new SMART. To assess the potential benefits in terms of participants' welfare gained by implementing a SMART-EXAM compared to a conventional SMART, we simulate data from these designs based on the original data. Note that for the illustration purpose, we assume the intermediate decision point is at one fixed time point for the SMART ADHD study thus ignoring the variable O_{21} . The specification for the required parameters is provided in Web Table 4 in Web Appendix D.

We assume equal treatment capacities, i.e., C = 0.5 for the SMART and SMART-EXAMs. For SMART-EXAMs, we specify $\eta = -1$, m = 1, and consider various values of $\epsilon = 0.1, 0.2, 0.3$. As stated in Section 3.1, even though the original SMART ADHD study did not collect participants' preferences for stage-2 treatments, it is reasonable to assume that, participants with lower adherence to the initial treatment, i.e., $O_{22} = 0$, are more likely to prefer $A_2 = 1$, i.e., augmentation, than those with higher adherence, i.e., $O_{22} = 1$. This corresponds to the scenario with a positive association between treatment effects and preferences in Section 4.1. We consider three possible settings for the preference data based on the above assumption, which is detailed in Web Table 4 in Web Appendix D.

5.1 Application results

As shown in Table 2, the SMART-EXAM with $\epsilon = 0.1$ performs worse than the other designs in terms of selecting the true optimal DTR, denoted by $\Pr(\hat{d}^{*IPW} = d^*)$, in all these scenarios. However, it outperforms the other designs in terms of improving the probability of receiving the preferred treatment and the mean outcomes. As ϵ increases, the learning ability of the SMART-EXAM improves and becomes closer to that of the SMART, while still continuing to improve $E[\bar{u}]$ and $E[\bar{Y}]$ compared to the SMART. These findings convey that, based on the original data, when conducting a new SMART for children with ADHD, the SMART-EXAM can increase the chance of allocating participants to their preferred treatment and improve the outcomes for enrolled participants, while still maintaining its effectiveness in estimating optimal DTRs with a moderate value of ϵ .

[Table 2 about here.]

6 Discussion

A conventional SMART randomizes participants with non-individualized and balanced randomization probabilities. Despite its implementation simplicity and desirable performance in comparing embedded DTRs, it faces some ethical issues. To mitigate these issues, we propose the SMART-EXAM, a novel SMART design that incorporates participants' preferences and predicted treatment effects into the randomization procedure, to advance health promotion among both the participants enrolled in the trial and future patients. We provide a detailed illustration of conducting a SMART-EXAM and assess its empirical properties through an extensive simulation study. The simulation results demonstrate that the SMART-EXAM can improve the average probability of receiving the preferred treatment and the final outcome compared with the conventional SMART, while also achieving a comparable ability to construct optimal DTRs for future patients with a moderate value of ϵ , which controls the range of randomization probabilities. We argue that by incorporating participant preferences and predicted treatment effects, the SMART-EXAM has the potential to encourage more people to participate in the trial and alleviate the lost-to-follow-up issue, finally resulting in a more diverse and representative sample compared to the conventional SMART.

It is important to note that each SMART design has its own strengths and weaknesses. The choice of design should depend on the primary research goals and specific contexts. For instance, if patient enrollment is challenging due to their reluctance to get randomized with balanced probabilities, the SMART-EXAM may be a better choice. Conversely, if the primary focus is on improving treatment effectiveness for enrolled participants, the SMART-AR may be more appropriate, as it continuously updates randomization probabilities based on accumulating information about treatment effects. When maximizing statistical power becomes a priority, the conventional SMART may be a suitable choice. The SMART-EXAM offers investigators the flexibility to adjust experimental parameters based on their specific research focus and clinical experience. Specifically, when the primary focus is to maximize the statistical power for future patients instead of participants' welfare, the SMART-EXAM can be *reduced* to a SMART with non-individualized randomizations by setting $\Lambda_i = \Lambda_j$ and $\zeta_i = \zeta_j, j \neq i$ for all participants. In addition, higher values of ϵ can also increase the learning ability of the SMART-EXAM. In contrast to many adaptive trials that lack strict capacity constraints, the SMART-EXAM incorporates a pre-specified treatment capacity $C_{a_2|a_1}$ to control the expected number of participants in each treatment. This feature resembles the SMART where the sample size is predetermined for each treatment arm and is particularly useful when treatment resources are limited or when investigators wish to prespecify the sample size for each treatment for the purpose of grant application.

The SMART-EXAM can be potentially extended to SMARTs with more than two stages or more than two treatment options at each stage. Furthermore, Walter et al. (2017) summarized a statistical framework for estimating treatment effects, selection effects, and preference effects based on data from patient-preference designs introduced in Section 1. The SMART-EXAM serves as another potential design to quantify these estimands.

Despite its great potential, several limitations of the SMART-EXAM need to be acknowledged. One limitation of the SMART-EXAM, which is also shared by various types of adaptive designs (Robertson et al., 2023), is that there are no explicit sample size calculation formulae for these designs. Simulation studies based on postulated parameters are needed to ensure sufficient statistical power. Given a sufficient sample size, the SMART-EXAM is expected to effectively identify the optimal DTR due to the consistency of the IPW estimators of DTR means, i.e., as the sample size increases, the estimated DTR means all converge to their true values, enabling the accurate identification of the true optimal DTR. Additionally, participants in a SMART-EXAM may feel disappointed when receiving a treatment they dislike after they have indicated their preferences. To manage expectations and minimize potential disappointment, it is crucial to ensure that participants have a realistic understanding of the SMART-EXAM by emphasizing that while this design could improve the average probability of receiving the preferred treatment, it cannot guarantee that every participant will be allocated to their preferred treatment. Another potential solution is to predict the preference using previous data resources instead of letting participants indicate their preferences. An additional challenge in a SMART-EXAM is the existence of participants without specific treatment preferences. One potential approach is to use balanced probabilities for participants without preferences, a straightforward adjustment that can be integrated into the existing SMART-EXAM algorithm. In addition, researchers may consider combining the no-preference group with one of the preference groups based on the specific context. For example, Chakraborty et al. (2013) merged participants who had no preferences with those who preferred augmentation with a new treatment. This combined group can be considered as participants open to either treatment option.

ACKNOWLEDGMENTS

We would like to thank Haruka Kiyohara, the referees, the associate editor, and the co-editor for providing insightful comments and suggestions to improve the paper. B. Chakraborty would like to acknowledge support from an Academic Research Fund Tier 2 grant (MOE-T2EP20122-0013) from the Ministry of Education, Singapore, as well as the start-up grant from the Duke-NUS Medical School, Singapore.

DATA AVAILABILITY

The data used in this paper is publicly available at http://www-personal.umich.edu/ ~dalmiral/atsworkshops.html.

References

- Almirall, D., Nahum-Shani, I. B., and Murphy, S. A. (2011). ADHD data. http: //www-personal.umich.edu/~dalmiral/atsworkshops.html, Accessed on September 6, 2023.
- Chakraborty, B., Laber, E. B., and Zhao, Y. (2013). Inference for optimal dynamic treatment regimes using an adaptive m-out-of-n bootstrap scheme. *Biometrics* **69**, 714–723.
- Cheung, Y. K., Chakraborty, B., and Davidson, K. W. (2015). Sequential multiple assignment randomized trial (SMART) with adaptive randomization for quality improvement in depression treatment program. *Biometrics* **71**, 450–459.
- Dumville, J., Hahn, S., Miles, J., and Torgerson, D. (2006). The use of unequal randomisation ratios in clinical trials: a review. *Contemporary Clinical Trials* **27**, 1–12.
- Edwards, S. J. and Braunholtz, D. A. (2000). Can unequal be more fair? a response to Andrew Avins. *Journal of Medical Ethics* **26**, 179–182.

Eisele, J. R. (1994). The doubly adaptive biased coin design for sequential clinical trials.

Journal of Statistical Planning and Inference 38, 249–261.

- Fava, M., Rush, A. J., Trivedi, M. H., Nierenberg, A. A., Thase, M. E., Sackeim, H. A., Quitkin, F. M., Wisniewski, S., Lavori, P. W., Rosenbaum, J. F., et al. (2003). Background and rationale for the sequenced treatment alternatives to relieve depression (STAR * D) study. *Psychiatric Clinics of North America* 26, 457–494.
- Fukami, T. (2023). Shared decision making with psychological safety. The Lancet 401, 1153–1154.
- Gaines, B. J. and Kuklinski, J. H. (2011). Experimental estimation of heterogeneous treatment effects related to self-selection. American Journal of Political Science 55, 724–736.
- Gunlicks-Stoessel, M., Mufson, L., Westervelt, A., Almirall, D., and Murphy, S. (2016). A pilot SMART for developing an adaptive treatment strategy for adolescent depression. *Journal of Clinical Child & Adolescent Psychology* 45, 480–494.
- Johnson, G. S., Topp, A. S., and Wahed, A. S. (2021). Optimizing DTRs via threshold utility analysis on quality adjusted survival. *Journal of Statistical Research* 55, 127–145.
- Knaus, M. C., Lechner, M., and Strittmatter, A. (2021). Machine learning estimation of heterogeneous causal effects: Empirical monte carlo evidence. *The Econometrics Journal* 24, 134–161.
- Knox, D., Yamamoto, T., Baum, M. A., and Berinsky, A. J. (2019). Design, identification, and sensitivity analysis for patient preference trials. *Journal of the American Statistical Association* **114**, 1532–1546.
- McKay, J. R., Drapkin, M. L., Van Horn, D. H., Lynch, K. G., Oslin, D. W., DePhilippis,
 D., Ivey, M., and Cacciola, J. S. (2015). Effect of patient choice in an adaptive sequential randomization trial of treatment for alcohol and cocaine dependence. *Journal of Consulting and Clinical Psychology* 83, 1021.

- Murphy, S. A. (2005). An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine* **24**, 1455–1481.
- Nahum-Shani, I., Qian, M., Almirall, D., Pelham, W. E., Gnagy, B., Fabiano, G. A., Waxmonsky, J. G., Yu, J., and Murphy, S. A. (2012). Q-learning: a data analysis method for constructing adaptive interventions. *Psychological Methods* 17, 478–494.
- Narita, Y. (2021). Incorporating ethics and welfare into randomized experiments. Proceedings of the National Academy of Sciences 118, e2008740118.
- Oetting, A. I., Levy, J. A., Weiss, R. D., and Murphy, S. A. (2011). Statistical methodology for a SMART design in the development of adaptive treatment strategies. In Shrout, P., Keyes, K., and Ornstein, K., editors, *Causality and Psychopathology: Finding the Determinants of Disorders and their Cures*, pages 179–205. Oxford University Press, New York.
- Pelham Jr, W. E., Fabiano, G. A., Waxmonsky, J. G., Greiner, A. R., Gnagy, E. M., Pelham III, W. E., Coxe, S., Verley, J., Bhatia, I., Hart, K., et al. (2016). Treatment sequencing for childhood ADHD: A multiple-randomization study of adaptive medication and behavioral interventions. *Journal of Clinical Child & Adolescent Psychology* 45, 396– 415.
- Robertson, D. S., Lee, K. M., López-Kolkovska, B. C., and Villar, S. S. (2023). Responseadaptive randomization in clinical trials: From myths to practical considerations. *Statistical Science* 38, 185 – 208.
- Rosenberger, W. F. and Hu, F. (2004). Maximizing power and minimizing treatment failures in clinical trials. *Clinical Trials* 1, 141–147.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* **66**, 688–701.

Rücker, G. (1989). A two-stage trial design for testing treatment, self-selection and treatment

preference effects. Statistics in Medicine 8, 477–485.

- Torgerson, D. J., Klaber-Moffett, J., and Russell, I. T. (1996). Patient preferences in randomised trials: threat or opportunity? *Journal of Health Services Research & Policy* 1, 194–197.
- Walter, S., Turner, R., Macaskill, P., McCaffery, K., and Irwig, L. (2017). Beyond the treatment effect: evaluating the effects of patient preferences in randomised trials. *Statistical Methods in Medical Research* 26, 489–507.
- Wang, J., Wu, L., and Wahed, A. S. (2022). Adaptive randomization in a two-stage sequential multiple assignment randomized trial. *Biostatistics* 23, 1182–1199.
- Wang, L., Rotnitzky, A., Lin, X., Millikan, R. E., and Thall, P. F. (2012). Evaluation of viable dynamic treatment regimes in a sequentially randomized trial of advanced prostate cancer. Journal of the American Statistical Association 107, 493–508.
- Wang, X. and Chakraborty, B. (2023). The sequential multiple assignment randomized trial for controlling infectious diseases: A review of recent developments. *American Journal* of Public Health 113, 49–59.
- Wu, L., Wang, J., and Wahed, A. S. (2023). Interim monitoring in sequential multiple assignment randomized trials. *Biometrics* 79, 368–380.
- Xu, Y., Greene, T. H., Bress, A. P., Sauer, B. C., Bellows, B. K., Zhang, Y., Weintraub, W. S., Moran, A. E., and Shen, J. (2022). Estimating the optimal individualized treatment rule from a cost-effectiveness perspective. *Biometrics* 78, 337–351.

SUPPLEMENTARY MATERIALS

Web Appendices referenced in sections 3, 4, 5, and source code are available with this paper at the Biometrics website on Oxford Academic. R Code for the simulation study is available at https://github.com/cherylwaal/SMART-EXAM.



Figure 1. An example of a two-stage SMART design for children with ADHD. R denotes randomization. BMOD and MEDS represent behavioral modification and oral methamphetamine respectively.



Figure 2. The schematic plot for a SMART-EXAM that individualizes the stage-2 randomization probabilities to account for participants' welfare.



Figure 3. Simulation results for the setting with a pilot SMART of size $N^p = 200$. A, B, C, and D correspond to the sample size N = 200, 300, 400, 500 respectively. The left (right) panel of each subfigure corresponds to the negative (positive) association between preferences and treatment effects. $E[\mu_{\hat{d}^*}]$ denotes the Monte Carlo estimate of the true value of the estimated optimal DTR; $E[\bar{u}]$ denotes the Monte Carlo estimate of the probability of being assigned to the preferred treatment, and $E[\bar{Y}]$ denotes the Monte Carlo estimate of the mean outcome.

Table 1

Parameter specification for generating data from SMART designs in the simulation study. $\zeta_{2,1,i}$ is the treatment effect of $a_2 = 1$ compared to $a_2 = -1$ on the outcome; ϵ is the parameter that controls the range of the randomization probabilities generated by SMART-EXAMs; η is the coefficient for the treatment effect in the price function.

Variable	Specification
The stage-1 treatment A_1	$(A_1+1)/2 \sim \text{Bernoulli}(0.5)$
Intermediate response status ${\cal R}$	$R \sim \text{Bernoulli}(0.5)$
Tailoring variables O_{21} and O_{22}	Standard normal distribution $N(0,1)$
The outcome Y	For non-responders: $Y_i = 2 - A_{1i} + A_{2i} + 0.5A_{1i}A_{2i} - 0.5O_{21,i}A_{2i} + 0.5O_{22,i}A_{2i} + \tau, \tau \sim N(0, 3^2)$
	For responders: $Y_i = 3 + A_{1i} + \tau$, $\tau \sim N(0, 3^2)$
Preference indicator Λ_i	Negative association: $Pr(\Lambda_i = 1) = logit^{-1}(-0.2\zeta_{2,1,i}+1)$
	Positive association: $Pr(\Lambda_i = 1) = logit^{-1}(0.2\zeta_{2,1,i} + 0.5)$
ϵ and η	1) $\epsilon = 0.1, \eta = -1; 2$) $\epsilon = 0.2, \eta = -1; 3$) $\epsilon = 0.3, \eta = -1; 4$) $\epsilon = 0.2, \eta = -0.7; 5$) $\epsilon = 0.2, \eta = -0.4; 6$) $\epsilon = 0.2, \eta = -0.1$

Table 2

Table 2 The application results. $Pr(\hat{d}^{*IPW} = d^*)$ is the probability of selecting the true optimal DTR by IPW; $E[\mu_{\hat{d}^{*IPW}}]$ is the Monte Carlo estimate of the true value of the estimated optimal DTR by IPW; $E[\bar{u}]$ is the Monte Carlo estimate of the probability of receiving the preferred treatment; $E[\bar{Y}]$ is the Monte Carlo estimate of the mean outcome among non-responders. The true optimal DTR and the true DTR means are approximated using a simulated SMART of sample size 10,000.

Scenario	Design	$\Pr(\hat{\boldsymbol{d}}^{*IPW} = \boldsymbol{d}^*)$	$E[\mu_{\hat{d}^{*IPW}}]$	$E[\bar{u}]$	$E[\bar{Y}]$
S1	SMART	0.984	3.504	0.500	3.002
	SMART-EXAM $\epsilon=0.1$	0.922	3.468	0.817	3.263
	SMART-EXAM $\epsilon=0.2$	0.978	3.500	0.741	3.200
	SMART-EXAM $\epsilon=0.3$	0.984	3.504	0.664	3.136
S2	SMART	0.984	3.504	0.500	3.002
	SMART-EXAM $\epsilon=0.1$	0.942	3.480	0.792	3.156
	SMART-EXAM $\epsilon=0.2$	0.960	3.491	0.722	3.115
	SMART-EXAM $\epsilon=0.3$	0.980	3.502	0.652	3.074
S3	SMART	0.984	3.504	0.500	3.002
	SMART-EXAM $\epsilon=0.1$	0.940	3.479	0.745	3.245
	SMART-EXAM $\epsilon=0.2$	0.980	3.501	0.687	3.186
	SMART-EXAM $\epsilon=0.3$	0.988	3.506	0.629	3.128