Real-time Adaptive Randomization of Clinical Trials

November 3, 2024

Prof. Gui Liberali,[a,c] Prof. Eric Boersma,[b] Hester Lingsma,[c] Jasper Brugts,[b] Prof. Diederik Dippel,[d] Prof. Jan Tijssen[e], Prof. John Hauser[f]

a. Erasmus University, Rotterdam School of Management, Rotterdam
b. Erasmus MC, University Medical Center, Department of Cardiology, Rotterdam
c. Erasmus MC, University Medical Center, Department of Public Health, Rotterdam
d. Erasmus MC, University Medical Center, Department of Neurology, Rotterdam
e. Amsterdam UMC, University of Amsterdam, Department of Cardiology, Amsterdam
f. Massachusetts Institute of Technology (MIT), MIT Sloan School of Management, Cambridge, U.S.

Contact author: Gui Liberali, liberali@rsm.nl. Address for correspondence: P.O. box 1738. 3000 DR Rotterdam. Netherlands.

Acknowledgements: The authors are very thankful for comments and support from Eric Langhamer, Diana Toli, Gaelle Saint Hilary, and Laura Velasco, and for the clinical trial data from Servier and Duke University School of Medicine

Real-time Adaptive Randomization of Clinical Trials

Abstract

Objective: To evaluate real-time (day-to-day) adaptation of randomized controlled clinical trials (RCTs) with delayed endpoints – a "forward-looking optimal-experimentation" form of response-adaptive randomization (RAR). To identify the implied tradeoffs between lowered mortality, confidence intervals, statistical power, potential arm misidentification, and endpoint-rate change during the trial.

Study Design and Setting: Using data from RCTs in acute myocardial infarction (30,732 patients in GUSTO-1) and coronary heart disease (12,218 patients in EUROPA), we resample treatment-arm assignments and expected endpoints to simulate (1) real-time assignment, (2) forward-looking assignments adapted after observing a fixed number of patients ("blocks"), and (3) a variant that balances RCT and real-time assignments. Blinded RTARs adjust day-to-day arm assignments by optimizing the tradeoff between assigning the (likely) best treatment and learning about endpoint rates for future assignments.

Results: Despite delays in endpoints, real-time assignment quickly learns which arm is superior. In the simulations, by the end of the trials, real-time assignment allocated more patients to the superior arm and fewer patients to the inferior arm(s) resulting in fewer mortalities over the course of the trial. Endpoint rates and odds ratios were well within (resampling) confidence intervals of the RCTs, but with tighter confidence intervals on the superior arm and less-tight confidence intervals on the inferior arm(s) and the odds ratios. The variant and patient-block-based adaptation each provide intermediate levels of benefits and costs. When endpoint rates change within a trial, real-time assignment improves estimation of the end-of-trial superior-arm endpoint rates, but exaggerates differences relative to inferior arms. Unlike most RARs, real-time assignment automatically adjusts to reduce biases when real changes are larger.

Conclusion: Real-time assignment improves patient outcomes within the trial and narrows the confidence interval for the superior arm. Benefits are balanced with wider confidence intervals on inferior arms and odds ratios. Forward-looking variants provide intermediate benefits and costs. In no simulations, was an inferior arm identified as statistically superior.

Keywords: adaptive clinical trials, multi-armed bandits, response-adaptive randomization, temporal changes, patient beneficence, forward-looking trials.

Running title: Real-time adaptive randomization. Text word count: 4,514

Plain language summary. Randomized control tests (RCT) are the gold standard in clinical trials – typically half of the patients are assigned to a new drug or procedure and the other half to a placebo (or the current best option). Typically, half of the patients might get an inferior drug or treatment. We explore a method, Real-time Adaptive Randomization (RTAR), that uses information observed up to the time of the next assignment to best allocate patients to treatments—balancing known current and unknown future outcomes—treating versus learning. RTAR is based on a preplanned, but adaptive, assignment rule. Blinding can be maintained, so that neither the trialist nor the patient knows to which treatment the patient was assigned. During the trial, as the RTAR learns the "best" treatment, the RTAR assigns more patients to that best treatment than would a classical RCT. In two large-scale cardiovascular clinical trials, our simulations suggest that the RTAR would have saved lives while identifying the best post-trial treatment at least as well as an RCT. Some statistical measures are improved and others are worse. If endpoint rates in treatments would have changed dramatically during the trial, the RTAR would have adapted better than many other methods.

Real-time Adaptive Randomization of Clinical Trials

1. Adaptive trials, response-adaptive-randomization, and real-time adaptive randomization

Relative to randomized controlled trials (RCT) patient lives might be saved (and non-fatal endpoints prevented) if blinded information, gained from patients within a large-scale trial, is used to automatically reassign more patients to the superior arm. (Superiority is identified during the trial.) We examine whether such increased patient benefit (beneficence) and assigning the best arm if known (equipoise) comes at the cost of higher-variance odds-ratios, a change in the ability to identify the best arm, or statistical confidence that the superior arm is indeed best.

Automatic, real-time adaptive randomization is a response-adaptive randomization (RAR) method [1, 2, 3, 4, 5, 6], which are themselves a type of adaptive designs such as when Data Safety Monitor Boards (DSMBs) periodically review results and reallocate the next batch of patients among arms [7, 8]. RAR methods include Thompson Sampling (assign patients proportional to the probability that an arm is best), modifications of Thompson sampling, play the winner, sequential maximum likelihood, sequential posterior mean, and various other methods based on Bayesian updating [2, 3, 4, 5, 6]. RARs vary in how they choose to use information in their adaptive-sampling strategies, leading to different trade-offs between patient beneficence and uncertainty reduction (e.g., statistical power for endpoint rates or for odds ratios). Because RARs often allocate more patients to the superior arm as the trial progresses, RARs may or may not be robust if (true) endpoint rates change during the trial (temporal changes) [9, 10, 11]. Furthermore, RAR analysis must use all available information and account for any small-sample biases that might be due to adaptivity [8, 10, 11, 12, 13, 14, 15]. When RARs are Bayesian in nature, reported statistics must be justified as appropriate for the data-

generating process [10, 16, 17].

RARs tend to be myopic, require randomization within groups of assigned patients, and use fixed overall sample sizes [12]. Recently proposed forward-looking patient assignment balances the benefit of learning the endpoint rates of the arms to make better patient assignments during the remainder of trial (and post-trial) against the immediate expected best-arm assignments in the current period [9, 13, 18, 19, 20].

Forward-looking methods show promise for increased beneficence and equipoise relative to myopic RARs and fixed randomization [19], but forward-looking optimization faces theoretical challenges when endpoints are delayed. Randomizations assume a trialist observes prior endpoints before randomization rates are changed. However, in our first empirical example, mortality is observed 30 days after arm assignment and in our second empirical example the last primary endpoints are observed with a mean of 4.2 years after treatment. To address delayed outcomes and to maintain randomization for every assignment, most previously-proposed forward-looking algorithms group patients into blocks of patients and sample from all potential patient orders within a block. Randomization rates are changed block to block [19].

Assigning patients within blocks is a creative and effective strategy, but does not fully address delayed endpoints. Unless the trialist plans no-assignment periods between successive blocks, delayed patient endpoints for assignments late in the block period are not observed in time for the next-block patient assignments. Large blocks reduce the percentage of unobserved outcomes, but large blocks decrease the advantages of optimal experimentation [19].

We examine an alternative forward-looking optimal-experimentation algorithm which assigns patients on a real-time (day-to-day) basis based on all data observed up to the day of patient assignment – real-time adaptive randomization (RTAR). Our analyses suggest that, had

RTAR assignments been used in two large-scale cardiovascular trials, lives would have been saved and non-fatal cardiovascular events prevented. We evaluate advantages, disadvantages, and ethical issues of RTARs. In §4.2, we examine the impact of temporal changes in endpoint rates, a known issue with RARs [9, 10, 11].

1.1. Multi-arm bandit algorithm when there are no delayed outcomes

RTARs use a preplanned statistical algorithm to assign patients to arms based on the endpoints observed up to the time of assignment. For ease of exposition, we first summarize a real-time adaptive design in which exactly one patient arrives each day and endpoints are observed the day of assignment. We next extend the discussion to the more-realistic situation where more than one patient arrives each day and endpoints are delayed. Based on patient homogeneity and tests that the true endpoint rates (e.g., mortality) per arm do not change throughout the trial (stationarity, §4.1), we assume that patients are interchangeable.

RTARs are based on multi-arm bandits (MAB) [21] where the trialist seeks to optimize endpoints over all current and future patients, including those after the trial. To best assign arms to patients, the trialist balances learning about the endpoint rates ("learning") and assigning the treatment most likely to be best ("treating"). See Figure 1. RTARs use Bayesian thinking – updating posterior beliefs about the distributions of endpoint-outcome probabilities after each day's (possibly delayed) outcomes.

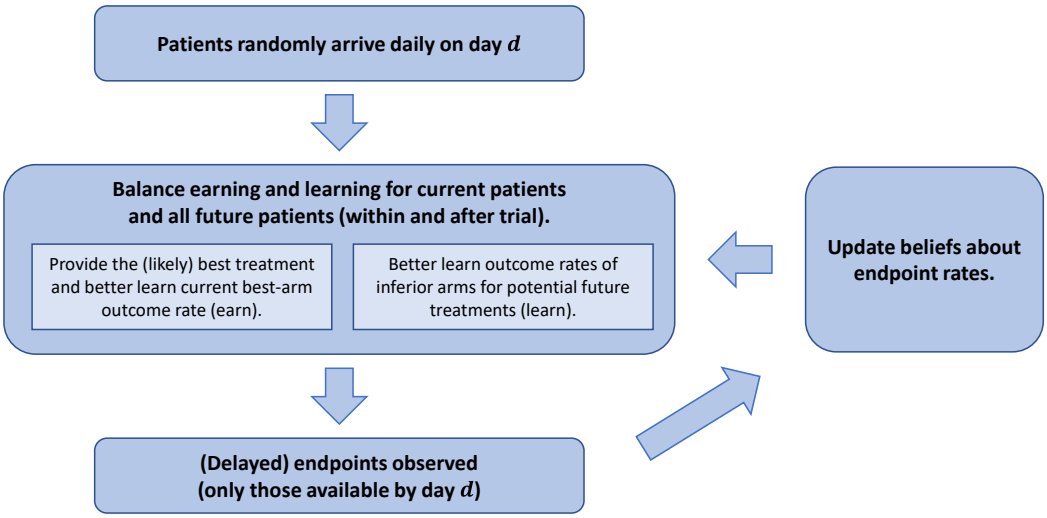

Fig. 1. Real-time adaptive randomization

If we know the best treatment today rather than later, we can save lives and prevent non-fatal cardiovascular events while we are learning. To capture this concept, trialists use a "discount" parameter to value endpoints today slightly more than endpoints tomorrow [13, 19, 21, 22, 23]. The "discount" parameter is chosen conservatively; analyses for our data suggest RTAR benefits for other (larger and smaller) discount-parameters values. We seek an algorithm that minimizes discounted negative endpoint outcomes for all future periods, including post trial.

eAppendix A describes full technical details about RTARs. We summarize the concepts here. An RTAR represents system knowledge about the endpoint rates using a Beta probability distribution per arm [9, 13, 24]. The Beta probability distribution has two parameters which are automatically updated when endpoints are observed. Based on this updating, an MAB-based RTAR chooses the patient assignments that best balance learning (potentially choosing suboptimal arms to learn more about them) and treating (always choosing the best arm given the available information).

Gittins proved that the optimal solution to the MAB is to compute a "Gittins index" that is a function of the updated parameters [13, 23]. The optimal policy is to assign to the arriving

patient the arm with the lowest Gittins index. (Lowest for mortality endpoints; highest for survival.) The calculations are easily completed and tabled before (blinded) patients are assigned to arms.

The RTAR algorithm differs from many RARs because patients are assigned to arms deterministically based on the endpoints observed up to that day. Because patients arrive randomly, randomization occurs over the trial. (Variants introduced in §1.3-1.4 allow randomization within days or blocks.) RTARs avoid arm assignments that are "unnecessary" for the learning process [13, 24]. RTARs automatically assign a sufficient number of patients to all study arms, but not necessarily in a 1:1 ratio, until uncertainty in endpoint rates is reduced enough that assigning patients to the inferior arms no longer provides value. In our experience, a trial size that is sufficiently powered as if run as an RCT is sufficiently long for real-time assignments to stabilize. In a new trial, trial size can be planned with simulation [8].

When patients arrive randomly and when endpoint rates do not change over the course of the trial, the random arrival of patients assures that the RTAR algorithm is an optimal randomization procedure. DSMBs still maintain independent oversight and monitoring of the trial progress, data integrity, and participant safety [8, 25].

1.2. Multi-arm bandit algorithm with delays and multiple patients per day

RTARs modify Gittins' solution. If more than one patient arrives on a given day, we assign all patients to the arm with the lowest index. When there are delays, we use only endpoints that have been observed by day d. The Gittins algorithm is no longer provably optimal, but we expect the algorithm to be close to optimal if (1) the number of patients that arrive on each day is small compared to the total patients in the trial and (2) the delay is small compared to the length of the trial. The first condition is met in both trials that we analyze, but the

performance of an RTAR remains an empirical question. The second condition is met in the first trial we analyze, but not necessarily the second trial, thus enabling us to examine the impact of substantial endpoint delays.

1.3. Block-based MAB based on Gittins' solution

To achieve randomization for every patient, a block-based MAB algorithm changes assignment probabilities only after a blocks of patients arrives and outcomes are observed [9, 19]. Empirically, if the block is sufficiently large relative to the endpoint delay, then most, but not all, endpoints can be observed before assignments are made in the next block. The algorithm samples the expected percentages of arm assignments over all possible patient orders. The block-based MAB is also known as the forward-looking Gittins index algorithm [FLGI, 9, 18, 19].

In simulations grounded to a breast-cancer-treatment RCT, the block-based MAB algorithm provided "substantial improvements in terms of patient benefit" relative to other trial strategies including RCTs and other RARs [19]. The block-based MAB improved the expected number of positive endpoints by almost 50%, but with a reduction in statistical power of approximately 70%. Results depended upon the block size, with more positive endpoints and lower power observed for smaller-sized blocks. Other RARs produced intermediate patient successes and power relative to the block-based MAB and an RCT.

1.4. RTAR η-variant to ensure a target minimum power

An RTAR assigns substantial sample to the superior arm, but less sample to inferior arms resulting in less statistical power for the inferior-arm endpoint rates. To address this ethical dilemma, trialists may wish to assure a minimum sample size (minimum statistical power) on the inferior arms or on odds ratios [9, 26].

An η–variant of an MAB algorithm addresses this dilemma by seeking a minimum level

of statistical power to inferior arms. With probability η, the η-variant randomizes patients in equal proportions to all arms that have not yet reached a targeted minimum number of patients. Otherwise, the η-variant assigns patients with Gittins' indices. The η-variant is an alternative means to achieve burn-in [3, 6, 27]. RTARs, the block-based MAB, and the η-variants are all MAB-based algorithms. An MAB-based algorithm is a type of RAR and an RAR is a type of adaptive design.

2. Statistical concepts, adaptivity bias, and expected performance

eAppendix B provides details. We provide summaries in this section.

2.1. Statistical concepts and potential adaptivity bias

- The likelihood function and the posterior distribution do not depend explicitly on how the RTAR assigns patients [16].
- All information about unequal sample sizes among arms is included in the likelihood function [10, 16, 17].
- Typically-used maximum-likelihood estimators (MLEs) can be reported and analyzed after the trial is completed, especially for large samples [12, 14, 15, 16, 28, 29].
- MLEs are consistent (asymptotically unbiased), but they may be biased for small samples [8, 11, 13, 14, 15, 16, 30, 31, 32]. Such biases are minimal for the large samples in the trials analyzed in this paper [15, 16, 33].
- When the number of patients is sufficiently large, resampling consistently generates all commonly-reported post-trial statistics [34].

2.2. Anticipated performance of an RTAR relative to an RCT

Because an RTAR usually allocates more patients to the superior arm, we expect:

- Fewer negative endpoints with an RTAR relative to an RCT.

- Tighter confidence intervals and more power for the superior arm, less-tight confidence intervals and lower power for inferior arms.
- For two arms, an RCT maximizes pairwise power and minimizes odds-ratio confidence intervals. For three or more arms, predictions are less clear.
- The η-variant approaches an RTAR as $\eta \to 0$ and approaches an RCT as $\eta \to 1$, enabling a trialist to fine-tune patient beneficence versus power for odds ratios.
- RTARs have advantages relative to other RARs for temporal changes in endpoint rates (§4.2).

3. What if the GUSTO-1 and EUROPA trials had been adapted in real time?

3.1. The GUSTO-1 and EUROPA RCT trials

To study the potential performance of RTARs, we use resampling simulations grounded by the data from the GUSTO-1 and the EUROPA trials [35, 36]. The design and principal results of both trials have been published and are summarized in eTable 1 in eAppendix C. Briefly, GUSTO-1 randomized a total of 31,180 patients presenting with acute myocardial infarction to one of three thrombolytic strategies. (30,732 patients after excluding observations with missing data. A fourth strategy was added later into the trial.) The primary endpoint was 30-day all-cause mortality and was lowest in the patients randomized to accelerated tissue plasminogen activator (t-PA) with intravenous heparin. The GUSTO-1 investigators concluded that this combination "is the best thrombolytic strategy to date (i.e., 1993) for patients with acute myocardial infarction."

The EUROPA investigators randomly assigned 12,218 patients with stable coronary heart disease to either a treatment with the angiotensin-converting-enzyme (ACE) inhibitor perindopril or to a matching placebo. The primary endpoint was a composite of cardiovascular death, non-fatal myocardial infarction and cardiac arrest with successful resuscitation, and was lowest in

those patients randomized to perindopril. In 2003, the investigators concluded that, "on top of other preventive medications, [perindopril] should be considered in all patients with coronary heart disease." The mean time-to-the-last observation of outcomes across all patients in EUROPA was 4.2 years after the start of the trial.

The GUSTO-1 and EUROPA trials were conducted according to the prevailing ethical regulations at the time, which included approval of the protocol by the institutional review board at the participating hospitals, and informed consent by the study participants. Our analyses are based on the individual (anonymized) patient data from the trials, which we obtained by courtesy of Duke University School of Medicine and Servier.

3.2. Data and grounded simulations

The detailed distribution of the RCT randomizations and endpoints per day in the GUSTO-1 and EUROPA trials are presented in online eAppendix D (eFigures 1 and 2). Data are displayed from the first randomization until the last observed primary endpoint.

Using the empirical trial data, we resampled patients to simulate what would have happened had the trial been based on an RTAR. Priors were weakly informative and equal for all arms, thus starting with an equally-likely ratio (1:1:1 for GUSTO-1; 1:1 for EUROPA). For each day of the trial, the RTAR automatically assigns patients arriving on day d to one of the study arms, based on observed endpoints up to that the beginning of day d. Patients for each arm are drawn randomly (with replacement, given stationarity and exchangeability of patients) from the pool of RCT patients in the chosen arm. To avoid a particularly favorable draw and to compute confidence intervals for all statistics, we repeat the process with 200 replicates for each study. In GUSTO-1, these pools have 10,255 patients in arm 1, 10,268 patients in arm 2, and 10,209 patients in arm 3. In EUROPA, these pools have 6,100 patients in Perindopril and 6,108 in the

placebo. The empirically-grounded simulations continue until the final day of the original RCT trials. Mean endpoint rates, confidence intervals, power, pairwise odds ratios, and other statistics of interest are based on the distributions over replicates.

3.3. Odds ratios

Throughout the trial, estimated odds ratios evolve and the odds-ratio confidence intervals become tighter. Figures 2a to 2c plot the evolution of the mean and the confidence intervals for the odds ratios of all pairs of arms (averaged over replicates). We observe a tighter confidence interval for the arm-1-to-arm-3 odds ratio (the two arms with lowest mortality rates) relative to the confidence intervals for the arm-1-to-arm-2 odds ratio (superior to third best) and for the arm-2-to-arm-3 odds ratio (second best to third best).

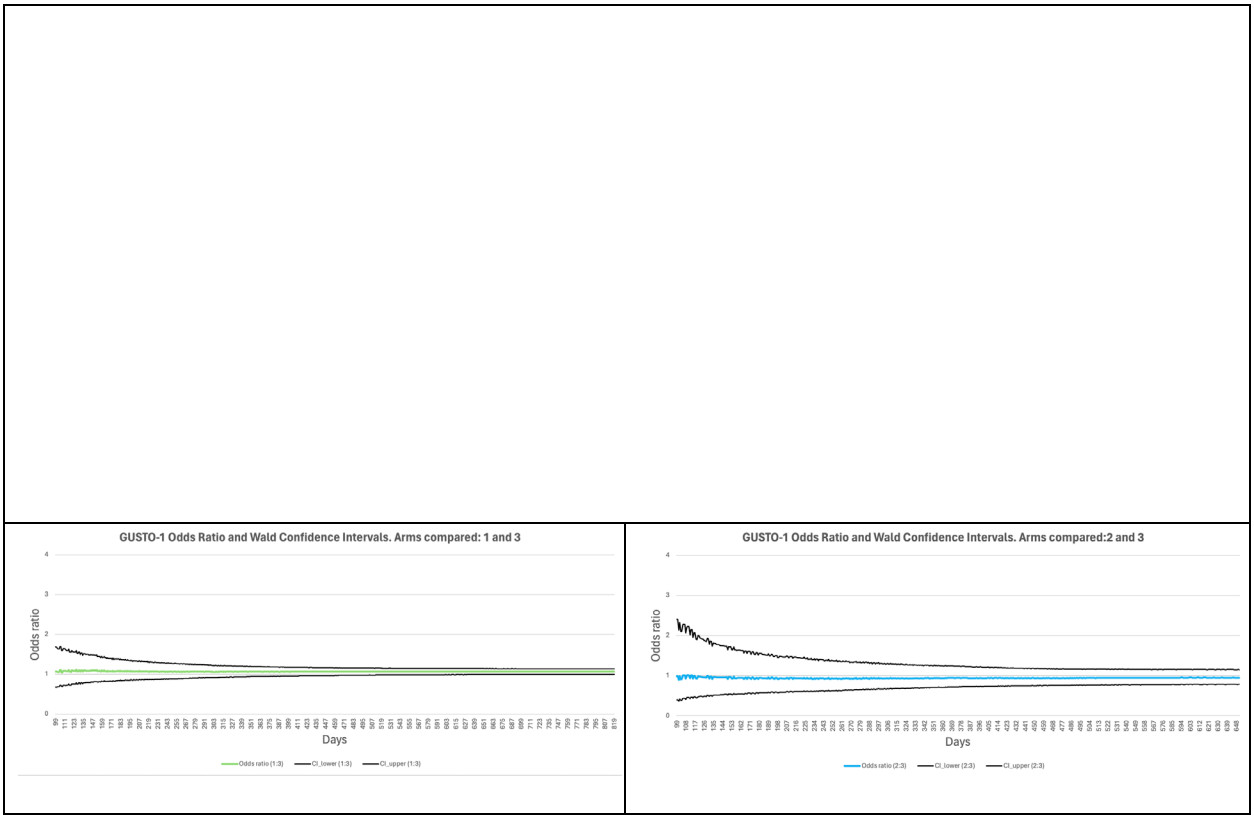

Fig. 2. Changes in RTAR odds ratios during the trial for arms 1 and 2 (top), arms 1 and 3 (bottom left), and arms 2 and 3 (bottom, right) averaged over 200 replicates

To obtain medians and confidence intervals, we ordered 2.5%, 50%, and 97.5% levels over 200 replicates. For GUSTO-1, the resampling median RCT arms 1:2 odds ratio is 1.182 (1.075, 1.297), the arms 1:3 odds ratio is 1.116 (0.997, 1.240), and the arms 2:3 odds ratio is 0.944 (0.852, 1.052). (In the RCT, the observed means were 1.184, 1.118, and 0.945, respectively.) The median RTAR odds ratio estimates are 1.204 (1.073, 1.632) for arms 1:2, 1.150 (1.032, 1.414) for arms 1:3, and 0.946 (0.721, 1.199) for arms 2:3. The RCT medians are within the RTAR confidence intervals and the RTAR medians are within the RCT confidence intervals for all pairs of arms. The confidence intervals for the RTAR odds ratios are wider than those for the RCT.

We get similar results for EUROPA. For example, the median odds ratio for Perindopril

versus a placebo is 1.246 (1.118, 1.412) and the median RTAR odds ratio is 1.248 (1.080, 1.455), $\eta = 0.25$.

3.4. Number of patients assigned, mortalities, and mortality rates

The first four columns of Table 1 present the results of the GUSTO-1 trial. The last three columns present the results of the EUROPA trial. We present the number of assigned patients, the number of primary endpoint events, and the endpoint rates for the original RCT (in the first three rows) along with confidence intervals. In the last nine rows of Table 1, we present the results had these trials used an RTAR, an η-variant, or a block-based MAB.

For both trials, the ranking of all arms in the simulations by the RTAR, the η-variant, and the block-based MAB match the RCT ranking (t-PA with IV Heparin is the best, SK with IV Heparin is the worst in GUSTO; Perindopril is the best, placebo is the worst in EUROPA). The primary endpoint rates estimated with all three adaptive algorithms are quite close to those estimated with the RCT and well within the confidence intervals. Relative to the RCT, all MAB variants provided tighter confidence intervals on the mortality rate for the (identified-within-the-trial) superior arm, with the tightest confidence interval provided for by the RTAR. As expected, the tighter bound for the superior arm comes with a tradeoff: confidence intervals are not as tight for the (identified-within-the-trial) inferior arms.

The lowest mortality (greatest beneficence) in GUSTO-1 was observed for the RTAR (1,952 lives lost) and the highest mortality for the RCT (2,074 lives lost) – a net saving of 122 lives due to real-time adaptation. The net savings for the η-variant and the block-based MAB were 72 and 102 lives saved, respectively.

Resampling suggests that the RCT would have identified the best arm in 98% of the replicates, comparable to the 99% achieved by the RTAR. There were no cases, for either the RTAR

or the RCT, where an inferior arm (arm 2 or arm 3) was identified as statistically significantly better than the superior arm (arm 1). There were only 2% cases of arm imbalance where one of the inferior arms (arm 2 or arm 3) was assigned more patients than the superior arm (arm 1). We obtain similar results for EUROPA despite the substantial delays in observing outcomes in the follow-up.

The gain in the reduction of negative endpoint outcomes comes at the cost of making fewer assignments to the inferior arms. For the RTAR, the assignments to the inferior arms averaged 2,755 (vs. 10,268) in GUSTO-1 arm 2, 4,311 (vs. 10,209) in GUSTO-1 arm 3, and 1,859 (vs. 6,108) in the EUROPA placebo arm. The η-variant and the block-based MAB allocated fewer patients to the superior arm and more patients to the inferior arms than the RTAR. For the η-variant, the observed minimums for inferior arms vary slightly from preset minimums because the real-time portion of the η-variant favors the superior arm.

Table 1. Benchmark simulation results for GUSTO-1 and EUROPA

	GUSTO-1				EUROPA		
	Arm 1[a]: t-PA, IV Heparin	Arm 2: SK, IV Heparin	Arm 3: t-PA+ SK, IV Heparin	Totals	Arm 1 [a]: Perindopril	Arm 2: Placebo	Totals
RCT					RCT		
Patients	10,255	10,268	10,209	30,732	6,110	6,108	12,218
Events	631	742	701	2,074	489	603	1,092
Event rate	0.062 (0.055, 0.068)	0.072 (0.068, 0.077)	0.069 (0.068, 0.069)		0.080 (0.074, 0.087)	0.099 (0.091, 0.106)	
Real-time adaptive randomization (RTAR)[b]					Real-time adaptive randomization (RTAR)[b, d]		
Patients	23,666	2,755	4,311	30,732	10,359	1,859	12,218
Events	1,455	200	297	1,952	828	183	1,012
Event rate	0.061 (0.059, 0.065)	0.073 (0.066, 0.096)	0.069 (0.064, 0.085)		0.080 (0.075 0.153)	0.098 (0.088, 0.227)	
Real-time adaptive randomization variant, $\eta = 0.25$ [b, c]					Real-time adaptive randomization variant, $\eta = 0.25$ [b,d]		
Patients	18,479	5,683	6,570	30,732	9,768	2,450	12,218
Events	1,136	413	453	2,002	782	241	1,023
Event rate	0.062 (0.059, 0.066)	0.073 (0.067,0.079)	0.069 (0.067, 0.076)		0.080 (0.075 0.086)	0.099 (0.088, 0.110)	
Block-based MAB [b, d]							
Patients	21,593	3,579	5,560	30,732			
Events	1,336	255	379	1,971			
Event rate	0.062 (0.059, 0.068)	0.071 (0.066, 0.081)	0.068 (0.066, 0.083)				

[a] Best arm in the trial

[b] Averaged over 200 replicates. Priors in GUSTO: $\alpha_o = 6$, $\beta_o = 390$ for all arms. Priors in EUROPA: $\alpha_o = 40$, $\beta_o = 1{,}800$ for both arms

[c] Minimum equally-likely allocation = 6,000.

[d] Block size = 60. Monte Carlo draws = 100.

3.5. The trade-off between (odds-ratio) statistical power and patient beneficence

More patients assigned to the superior arms implies greater power for the superior-arm' endpoint rates. Fewer patients assigned to inferior arms implies less power for inferior-arms' endpoint rates. The deviation from equal allocation of patients to arms implies lower power for the odds ratios [12, 18]. To examine this tradeoff further, we plot the change in odds-ratio statistical power (solid lines, left vertical axis) and the number of patient exposed to the superior and inferior arms (dotted lines, right vertical axis) for different values of η.

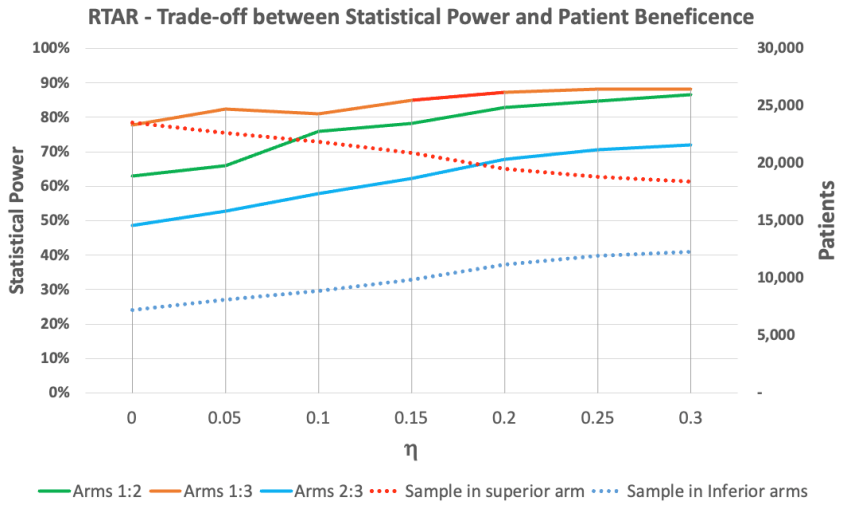

Fig. 3. Trade-off between statistical power and patient beneficence in GUSTO-1

Larger η's (more sample to 1:1:1 randomization) provide higher statistical power for the odds ratios, but lower patient beneficence (spread between the orange and blue dotted lines). At $\eta \geq 0.30$, the power of the η-variant is almost indistinguishable from the power of the RCT (90%). The trialist can choose η to make ethical tradeoffs between beneficence and odds-ratio power. (Open-source code provided.)

3.6. Greater patient beneficence with larger differences in endpoint rates

Every life is important, but it is beyond the scope of this article whether 122 fewer deaths out of 2,074 mortalities in GUSTO-1 justifies the use of a new method. However, if we examine

the GUSTO-1 trial, we see that the three arms are close in mortality risk, 0.062, 0.069, and 0.072. As an hypothetical, we examine more substantial differences—mortality rates of 0.063, 0.126, and 0.189 for the three arms, unknown before the trial. With 30,732 patients, an RTAR would have saved 1,700 lives compared to an RCT. We kept the total patients the same for a clear comparison. If the trialist had strong priors on the mortality risk and required the same statistical power, the trialist would allocate fewer patients to both the RTAR or the RCT. Even in this case, the RTAR would lead to substantially greater patient beneficence.

3.7. Summary of the GUSTO-1 and EUROPA empirically-grounded simulations

For GUSTO-1 and EUROPA, an RTAR increases patient beneficence, reduces patient risk, provides tighter confidence intervals and more power for the superior-arm endpoint rates, and comparable pairwise power for the superior-arm-to-second-best-arm comparison. An RTAR reduces power for inferior-arm endpoint rates and pairwise power for superior-to-third-best-arm and inferior-arm comparisons. The η-variant and the block-based MAB algorithm provide an intermediate balance of the benefits and costs of an RTAR versus an RCT. Odds-ratio power is comparable to an RCT for $\eta = 0.30$ or higher. Neither the RTAR, the η-variant, nor the block-based MAB identify an inferior arm as statistically superior.

4. Stationarity and temporal changes in endpoint rates

4.1. Tests of stationarity

To test whether endpoint rates change over the course of the trial, we split the RCT trial by quantiles on the date of assignment and examine whether endpoint rates vary significantly by quantile. For deciles, the null hypothesis of stationarity was not rejected for all GUSTO-1 arms (arm 1 $p = 0.47$, arm 2 $p = 0.45$, arm 3 $p = 0.80$) and for both EUROPA arms (Perindopril $p = 0.37$, placebo $p = 0.39$). Two-way quantile splits were also not significantly different. Other

stationarity tests are available from the authors.

4.2. Temporal changes in endpoint rates

Temporal changes in endpoint rates (non-stationarity), such as changes resulting from mutation in a virus, a change in the demographics of patients, environmental changes, or the advent of auxiliary treatments, are a known issue with RARs and potentially an issue with RTARs [9. 10. 11. 14]. eAppendix E explores the behavior of an RTAR with temporal shocks of 5%, 10%, 15%, 20% and 25% higher mortality for later patients than earlier patients. (The literature suggests that RARs are robust to drift rates less than 25%, but not for large rates [9. 14].)

As detailed in eTable 2 and consistent with the literature, RTAR estimates are closer than RCT estimates to "true" end-of-trial mortalities, but between-arm differences are upwardly biased. RTARs may have advantages relative to other RARs. For large shocks, an RTAR reexplores post-shock inferior arms resulting in reduced bias (only 1-3% for higher shock rates). This is an interesting topic for further exploration, especially in light of the developing theory for MABs that anticipate temporal changes [37, 38].

5. Discussion

Even when there are delays of 30 days or more, an RTAR would have saved lives (in GUSTO-1) and avoided cardiovascular events (in EUROPA) relative to an RCT while providing estimates of endpoint rates and odds ratios within statistical confidence of the RCT. The RTAR identifies well the superior arm for post-trial assignment. In no replicates was an inferior arm identified as statistically superior; the superior arm was identified as least as often with the RTAR as the RCT. Some statistical measures are better (tighter superior-arm confidence) and some worse (odds-ratio confidence intervals). Variants enable trialists to balance benefits and costs.

Ethically, RTARs enhance beneficence [2] and equipoise. However, patients who enter the trial late are more likely to receive the better treatments than patients who enter the trial early (violating equality).

GUSTO-1 and EUROPA were large-sample, stationary trials of medical *treatments*. Benefits might be less:

- For small-sample trials which require sufficient sample in inferior arms, such as many cancer trials, and reduced patient population trials such as rare-disease trials.
- When, during the course of the trial, there are temporal changes in endpoint rates or there are large changes in the treatment, such as technology changes in medical devices.
- When delays between randomization and endpoints are extremely large.

Endpoint rates were delayed more in EUROPA than GUSTO-1 suggesting that, while RTARs can handle delays, they might struggle with substantial delays. We encourage research on earlier adaptivity based on biomarkers.

Our simulations are empirically-grounded and their implications are as predicted by theory, but our simulations are *post hoc* analyses of the GUSTO-1 and EUROPA trials. There is nothing in our analyses that used knowledge that was not available at the time of RCT patient assignment. Nonetheless, any *post hoc* analyses must be treated with caution.

Our empirical results suggest that trialists can rely on random arrival with changing deterministic assignments rather than require an MAB to adjust randomized assignments after b patients have arrived ($b \gg 1$). For multiple-trial multiple-population settings, researchers can merge RTARs and platform-trials [35].

References

[1] Tricco, A, Tovey, D. Flexible approaches to clinical trials. Journal of Clinical Epidemiology 2023;154: A1-A2.

[2] Berry DA, Eick SG. Adaptive assignment versus balanced randomization in clinical trials: a decision analysis. Statistics in Medicine 1995;14:231-246.

[3] Du Y, Wang X, Lee JJ. Simulation study for evaluating the performance of response-adaptive randomization. Contemporary Clinical Trials 2015;40:15-25.

[4] Thall PF, Wathen JK. Practical Bayesian adaptive randomization in clinical trials. European Journal of Cancer 2007;43:589-866.

[5] Trippa L, Le EQ, Wen PW, Batchelor TT, Cloughesy T, Parmigaiani G, Alexander BM. Bayesian adaptive randomized trial design for patients with recurrent glioblastoma. Journal of Clinical Oncology 2012;30(2):3258-3263.

[6] Wathen JK, Thall PF. A simulation study of outcome adaptive randomization in multi-arm clinical trials. Clinical Trials 2017;14(5):432-440.

[7] Ryan, EG, Bruce J, Metcalfe AJ, Stallard N, Lamb SE, Viele K, Young D, Gates S. Using Bayesian adaptive designs to improve phase III trials: a respiratory care example. BMC Medical Research Methodology 2019;19(99):1-10.

[8] FDA. Adaptive designs for clinical trials of drugs and biologics: Guidance for industry. Biostatistics. Food and Drug Administration, Center for Drug Evaluation and Research, November; 2019; FDA-2018-D-3124.

[9] Villar SS, Bowden J and Wason J. Response-adaptive designs for binary responses: how to offer patient benefit while being robust to time trends? Pharm Stat 2018;17:182–19

[10] Proschan M, Evans S. Resist the temptation of response-adaptive randomization. Clinical Infectious Diseases 2020;71:3002-3004.

[11] Thall P, Fox P, Wathen J. Statistical controversies in clinical research: Scientific and ethical problems with adaptive randomization in comparative clinical trials. Annals of Oncology 2015;26:1621-1628.

[12] Hu F, Rosenberger WF. The theory of response-adaptive randomization in clinical trials. Wiley Series in Probability and Statistics. Wiley Interscience. Hoboken NJ. 2006.

[13] Villar SS, Bowden J, Wason J. Multi-armed bandit models for the optimal design of clinical trials: Benefits and challenges. Stat. Sci. 2015;30(2):199-215.

[14] Robertson DS, Lee KM, Lopez-Kolkovska BC, Villar SS. Response-adaptive randomization in clinical Trials: From myths to practical considerations. Statistical Science, 2023;38(2):185-208.

[15] Pallmann P, Bedding AW, Choodari-Oskooei B, Dimairo M, Flight L, Hampson LV, Holmes J, Mander AP, Odondi L, Sydes MR, Villar SS, Wason JMS, Weir CJ, Wheeler GM, Yap C, Jaki T. Adaptive designs in clinical trials: Why use them, and how to run and report them. MBC Medicine 2018;16(29):1-15.

[16] Bowden J, Trippa L. Unbiased estimation for response adaptive clinical trials. Statistical Methods in Medical Research 2017;26(5):2376-2388.

[17] Carlin BP, Louis TA Bayes and empirical Bayes methods for data analysis. Chapman & Hall/CRC, New York NY. 2000.

[18] Barnett HY, Villar SS, Geys H, Jaki T. A novel statistical test for treatment differences in clinical trials using a response-adaptive forward-looking Gittins Index Rule. Biometrics 2020;79:86-97.

[19] Villar SS, Wason J, Bowden J. Response-adaptive randomization for multi-arm clinical trials using the forward-looking Gittins Index rule. Biometrics 2015;71:969-978.

[20] Williamson SF, Villar SS. A response-adaptive randomization procedure for multi-armed clinical trials with normally distributed outcomes. Biometrics 2020;76:197–209.

[21] Berry D and Fristedt B. Bandit Problems – Sequential Allocation of Experiments; London: Chapman and Hall. 1985.

[22] Chick SE, Gans N, Yapar Ö. Bayesian sequential learning for clinical trials of multiple correlated medical interventions. Management Science 2022;68(7):4919-4938.

[23] Gittins JC 1979. Bandit processes and dynamic allocation indices. Journal of the Royal Statistical. Society 1979; (Ser B)41(2):148–177, plus commentary

[24] Gittins J, Glazebrook K and Weber R. Multi-armed bandit allocation indices. London: Wiley. 2011.

[25] Clemons F, Elbourne D, Darbyshire J, Pocock S, Damocles Group. Data monitoring in randomized controlled trials: Surveys of recent practice and policies. Clinical Trials 2005;2:22-33.

[26] Flehinger BJ, Louis TA, Robbins H, Singer BH. Reducing the number of inferior treatments in clinical trials, PNAS 1972;69(10):2993-2994.

[27] Granholm A, Kaas-Hansen BS, Lange T, Schjørring OL, Andersen LW, Perner A, Jensen AKG, Hylander Møller MH, An overview of methodological considerations regarding adaptive stopping, arm dropping, and randomization in clinical trials. Journal of Clinical Epidemiology 2023;153:45-54.

[28] Connor JT, Elm JJ, Broglio KR. Bayesian adaptive trials offer advantages in comparative effectiveness trials: an example in status epilepticus. Journal of Clinical Epidemiology 2013;66:S130-S137.

[29] Melfi VF, Page C. Estimation after adaptive allocation. Journal of Statistical Planning and Inference 2000;87(2):353-363.

[30] Pocock SJ. Group sequential methods in the design and analysis of clinical trials. Biometrika 1977;64(2):191-199.

[31] Pocock SJ (2006) Current controversies in data monitoring for clinical trials. Clinical Trials 2006;3:513-521.

[32] Todd S, Whitehead A, Stallard N, Whitehead J. Interim analyses and sequential designs in phase III studies. Clin. Phama. 2001;51:394-399.

[33] Hadad V, Hirshberg DA, Zhan R, Wager S, Athey S. Confidence intervals for policy evaluation in adaptive experiments. PNAS 2021;118(15):1-10.

[34] Efron B, Tibshirani RJ. An introduction to the bootstrap. New York NY :Chapman & Hall/CRC 1994.

[35] The GUSTO Investigators. An international randomized trial comparing four thrombolytic strategies for acute myocardial infarction. NEJM 1993;329(10): 673:682

[36] The European trial on reduction of cardiac events with perindopril in stable coronary artery disease investigators. Efficacy of perindopril in reduction of cardiovascular events among patients with stable coronary artery disease: randomized, double-blind, placebo-controlled, multicenter trial (the EUROPA study). Lancet. 2003; 362:782-788

[37] Chen Q, Golrezaei N, Bouneffouf N. Non-stationary bandits with auto-regressive temporal dependency. 37th Conference on Neural Information Processing Systems 2023:1-22.

[38] Liu Y, Xu K, van Roy B. Non-stationary bandit learning via predictive sampling. arXiv:2205.01970v6. 2023:1-42.