Statistical Science
2023, Vol. 38, No. 2, 212–215
https://doi.org/10.1214/23-STS865B
Main article: https://doi.org/10.1214/22-STS865

Comment: Response-Adaptive Randomization in Clinical Trials: From Myths to Practical Considerations

Yunshan Duan, Peter Müller and Yuan Ji

Robertson et al. [6] (henceforth RLLV) present a comprehensive review on existing response-adaptive randomization (RAR) approaches and related practical considerations. We congratulate the authors for a thorough and insightful discussion. The heart of the paper appears in Section 3 when the authors address some of the key questions about the use of RAR for clinical trials.

RAR adaptively modifies randomization probabilities based on the observed response data from patients during an ongoing clinical trial. Although randomization may be biased, it is important to recognize that the benefits of randomization are not entirely lost under an RAR approach. Randomization is a critical feature of clinical trial design as it allows causal attribution of differences in outcome to treatment assignment. This is achieved by balancing all potential confounders. The question then arises, to what extent might a RAR approach lose this ability to balance confounders? Most RAR approaches maintain some randomization (after all, "randomization" is part of the name). However, in the case of extremely inbalanced patient allocation, like the ECMO trial [1], the answer may be unclear. Fundamentally, why would one consider RAR instead of equal randomization (ER)? Statistical intuition might suggest balanced allocation of experimental units (e.g., ER) to be optimal, making biased allocation possibly suboptimal. However, as RLLV (Section 3.2, and Table 1) show in two-arm and multiarm trials with binary outcomes for some RAR designs, this is not always the case. Do the optimal ratios shown in Section 3.2 sacrifice other statistical properties (such as efficiency in inference or frequentist type I error rate) to achieve higher power than ER?

Technically, RAR uses observed outcomes to modify randomization probabilities. Since the observed data are associated with random noise, it is inevitable that the modification of the randomization probabilities may be "wrong," in the sense of increasing randomization probabilities for inferior arms. This is considered in Section 3.1. Does RAR lead to a substantial chance of allocating more patients to an inferior treatment? Putting aside the question about being substantial, in a finite sample size RAR clearly has nonzero probability of allocating more patients to an inferior treatment. The probability here is defined as (frequentist) probability under repeated use of RAR designs. As shown in Table 1 (RLLV), even under ER there is a 0.069 probability of allocating more than 10% of sample size to an inferior arm. Some RAR approaches avoid such imbalance by introducing explicit constraints. For example, not allowing imbalance to exceed 10% by construction implies zero probability of achieving more than 10% imbalance. Perhaps this is related to the reason for some of the zeroes (for imbalance beyond 10%) in Table 1 of RLLV. And, vice versa, some of the substantial imbalance metrics $\hat{S}_{0.1}$ for other RAR approaches in Table 1 are only a result of the lack of such constraints, rather than any fundamental differences in statistical modeling or algorithms across these RAR approaches. If desired, the same constraint could easily be added to any approach.

Like most literature, the authors focus the discussion of adaptive designs on the myopic problem of treatment assignment for the next patient (cohort). While most would agree that formally the problem is a sequential decision problem [2], Section 2.6, a full solution of the sequential decision problem is computationally infeasible and certainly not practicable in the setting of most clinical trial designs. However, some features of a sequential design could be retained. A common theme of optimal sequential decisions is the conflict of "exploration" versus "exploitation." In the early cohorts of a clinical trial, there might be a benefit for assigning some patients to possibly suboptimal doses or treatments because learning about a dose-response curve could help to eventually treat more patients at more effective doses or treatments. Of course, any such randomization could only be considered in a situation of approximate equipoise, or minor adverse events. In contrast, for the last few cohorts there is less benefit of learning, and one might want to focus more on the most promising doses or treatments. In other words, a design that starts with exploration and later only switches

Yunshan Duan is graduate student, Department of Statistics and Data Science, University of Texas, Austin, Texas 78705, USA (e-mail: yunshan@utexas.edu). Peter Müller is Professor, Department of Statistics and Data Science, University of Texas, Austin, Texas 78705, USA (e-mail: pmueller@math.utexas.edu). Yuan Ji is Professor, Department of Biostatistics, University of Chicago, Chicago, Illinois 60637, USA (e-mail: YJi@health.bsd.uchicago.edu).