Published in final edited form as:

J Am Stat Assoc. 2012 ; 107(498): 518–520. doi:10.1080/01621459.2012.665198.

Rejoinder to comments on Evaluation of Viable Dynamic Treatment Regimes in a Sequentially Randomized Trial of Advanced Prostate Cancer

Lu Wang, Andrea Rotnitzky, Xihong Lin, Randall E. Millikan, and Peter F. Thall

1 A BRIEF HISTORY OF THE SCRAMBLE

Because the analyses reported in this paper are the culmination of a process that began in 1998 when the trial was conceived by Millikan and Thall at M.D. Anderson Cancer Center (MDACC), it is worthwhile to first provide a brief history, and also some remarks on the relationship between theory and application.

The prostate cancer trial was motivated, in part, by the desire to better reflect actual clinical practice by evaluating multi-stage treatment strategies. This was a radical idea for oncologists in 1998 and, sadly, it remains so at this writing. This trial was dubbed “The Scramble” by members of the Genitourinary Medical Oncology Department at MDACC, due to its unique and apparently complex structure. When we conceived the prostate trial scramble design in 1998, we were not aware that a literature on DTRs existed. Since then, the scramble has served as a prototype for later oncology trials, including an ongoing trial at MDACC of six two-stage DTRs of targeted agents for advanced kidney cancer (Thall et al., 2007).

The immense research activity in DTRs among statisticians over the past decade seems to have been motivated in large part by their interesting mathematical structure and the challenging technical problems that they present. The original impetus for this area of research was analysis of complex observational data, pioneered by Robins (1986, 1987, 1989). Because DTRs are essentially mathematical formalisms of routine medical practice, clinical trials based on DTRs provide great advantages compared to conventional trials that reduce variables and focus on one stage of therapy. Evidently, as statisticians we have not done an effective job of communicating these facts to the oncology and larger medical research communities. In fact, actual trials of DTRs are few even as theoretical research activity has exploded. The good news is that many oncologists and research physicians in other areas very much like the idea of designing trials to evaluate DTRs, and moreover implementing such trials is not much harder than conventional trials. Since the ultimate aim of all this research activity is to develop improved medical therapies, it is high time for a more proactive approach to communicating these new statistical ideas to the medical research communities, and undertaking the hard work of actually implementing trials to evaluate DTRs.

2 SMARTS

We thank Almirall, Lizotte, and Murphy (ALM) for their kind words regarding our re-analysis of the data from the prostate cancer trial. ALM provide a nice overview of sequentially multiple assignment randomized trials (SMARTs), and focus on design and analysis as separate issues. ALM argue that it is important to design trials to evaluate viable DTRs by giving “a clear operationalization of how to treat patients,” for all foreseen outcomes, i.e. codifying the actions to be taken in response to all foreseen outcome histories.

We could not agree more. However, a viable DTR is one that actually can and will be implemented in the targeted population by the treating physicians. In most medical settings, there are a myriad of possible outcomes and accompanying actions. Inevitably, this complexity must be simplified in order to obtain a feasible design, including viable regimes that actually can and will be implemented in practice during actual trial conduct.

Stopping chemotherapy in patients with advanced disease when progressive disease (PD) or severe toxicity (TOX) occur is so routine in oncology practice that, when we conceived the prostate trial scramble design in 1998, it did not seem worth mentioning in the design. In fact, because all patients that developed PD or TOX were removed from the intervention, the removal rule was effectively operationalized in this trial. Years later, for this third data analysis, our inclusion of PD and TOX as outcomes and numerically scoring the resulting bivariate per-stage outcome were motivated, initially, by what we naively considered to be 47 dropouts. Formally and correctly calling PD and TOX outcomes led to 35 of these 47 patients being fully evaluated, and it enabled us to define viable DTRs that elaborated the simpler DTRs in the original design.

This trial raises three important points. The first is that one cannot optimally design an experiment until after it has been carried out. Only actually conducting a trial can reveal the many complexities that are relevant to the study goals. That is, science is a learning process. ALM provide a very useful explanation of this issue as it relates to SMARTs in their Section 2.2. The second point is that working harder to explain and define actual clinical outcomes led us to define DTRs that were viable, and hence that actually could be carried out in practice by treating physicians. The third point is that treatments have many effects, and accounting for these inevitably leads both to more complex statistical models and the need, as ALM point out, to account for trade-offs between efficacy and toxicity.

While this all may appear to be quite challenging, ALM point out more good news, namely that clinical trial protocols almost invariably include detailed explanations of the actual viable DTRs that will be used. After all, the first purpose of a clinical trial is to treat patients, and any reasonable protocol must spell out how this will be done. To simplify things somewhat, ideally, a medical statistician's task is to incorporate this sort of explanation formally into the trial design by defining the viable DTRs that will be evaluated. To motivate this, ALM provide a useful account of the consequences of having non-viable DTRs in a SMART trial. Their discussion of this issue could be used, in its simplest form, as an explanation of why conventional trial designs often fail to provide results that are useful to practicing physicians.

3 COMPOSITE ENDPOINTS

In their discussion of composite endpoints, ALM provide a simple and useful way to combine two of our scoring methods, by taking a weighted average. Because careful consideration of clinical events almost invariably leads to multivariate outcomes, inevitably, there arises the need to somehow reduce this to one dimension, in terms of either observations or parameters. ALM's repeated use of the words "trade-off" and "preference" in their discussion is encouraging, as we believe that this sort of thinking is essential when evaluating multiple, often competing clinical outcomes. For example, methods based on scores (Bekele and Thall, 2004), efficacy-toxicity trade-offs (Thall and Cook, 2004), and utilities (Houede et al., 2010) have been proposed for dose-finding trials, while quality adjusted survival time (Glasziou, et al. 1990; Zhao and Tsiatis, 1996) has been used for many years as a more honest and ethical tool than using survival time alone. These and similar methods should be portable for use in designing SMARTs, and we feel that they will play a central role when evaluating DTRs in the future.

4 NEW METHODS

We thank Chaffe and van der Laan (CVDL) for their kind remarks and their concise explanation of the formalism underlying our analyses. CVDL indicate that because in a SMART trial the treatment probabilities are known by design and bounded away from zero, IPW estimation of the DTR-specific counterfactual means offers an easy-to-implement procedure that yields estimators that are well behaved, i.e. nearly unbiased and with an approximately normal distribution. They reiterate the point we made in section 4.4, noting that IPW estimators can be made nearly efficient by replacing the known treatment probabilities by probabilities estimated under models for the treatment process that incorporate covariates that are predictive of the outcome. They go on to indicate that when, as in our trial, there are few drop-outs, IPW estimators that further weight by the inverse of the conditional probabilities of drop-out given past data offer an attractive practical analytic choice for SMART trials. We agree with all of these points. In fact, this is precisely the rationale that led us to analyze the data using the proposed IPW methodology.

Re-visiting arguments made in the past by Robins and colleagues (see e.g. Robins and Ritov, 1997) CVDL argue that, in the presence of severe drop-out, even if one is prepared to assume that drop-out is ignorable, i.e. that the drop-out probabilities are conditionally independent of the outcome given past covariates and treatments, due to the curse of dimensionality one cannot estimate the counterfactual DTR means without making some dimension reduction assumption about either (a) the drop-out probabilities or (b) the conditional law of the endpoint of interest given the data collected up to each study cycle. CVDL offer the TML methodology as an attractive analytic tool for SMART studies with severe ignorable drop-out. They indicate that TML is a general strategy for estimation of parameters of non parametric or semiparametric models. In the context of SMART clinical trials with ignorable drop-outs, certain variants of TML produce estimators of the DTR counterfactual means having the following attractive features:

1. They respect natural constraints on the counterfactual means,
2. they are double-robust, i.e they are consistent and asymptotically normal provided either a working model for (a) or a working model for (b) is correct, but not necessarily both,
3. they achieve the smallest possible variance for estimators that are consistent and asymptotically normal under the semiparametric model that assumes that the working model for (a) to be correct, if indeed the data generating law also satisfies the working model for (b), a property that CVDL refer to as local efficiency, and
4. they are guaranteed to be at least as efficient as any given IPW estimator that uses known or consistently estimated treatment and drop-out probabilities.

CVDL present a simulation study of a moderately sized SMART trial with two time points, no dropout and four possible static treatment regimes. In their simulation, TML performs better than an IPW estimator that uses treatment probabilities estimated under a logistic regression model with linear terms in the variables recorded up to each time point. Given that, in a clinical trial with no drop-outs, the treatment probabilities are known by design, double robustness is not an issue. For such settings, one could also use a simple, easy-to-implement variant of the IPW estimator which also satisfies (1), (3) and (4) (see, e.g. Section 6 of Rotnitzky and Robins, 1995). This is an IPW estimator in which the treatment probabilities are estimated under a logistic regression model that, in addition to the linear terms in the past covariates, includes specifically tailored pre-processed covariates computed at a first stage. It would have been interesting to include this estimator in CVDL's simulation study.

Several proposals other than TML exist that yield estimators satisfying some or all of (1)–(4). For example, Robins (2000) and Bang and Robins (2005) satisfy (2)–(3), Tan (2006, 2008) and Cao, Tsiatis and Davidian (2010) satisfy (2)–(4) and, Rotnitzky, Lei, Sued and Robins, (2012) and Tan (2010), satisfy (1)–(4). Whereas (2)–(4) are attractive theoretical large sample properties, it is important to bear in mind that many cancer studies are moderately sized, thus these properties are relevant only so long as they are good proxies for finite sample performance. We thus agree with the statement made by CVDL in their concluding remarks that more research is needed to understand the performance in terms of bias and efficiency of TML and any of the alternative semiparametric approaches in realistic SMART clinical trial settings.

5 CONCLUSIONS

Although it is well known that DTRs reflect more closely what happens in clinical practice, their use in clinical studies is still in its infancy, and there are many practical challenges. While increasing statistical methodological research has been done in recent years on design and analysis of DTRs, there are still many open methodological issues. Some of these were encountered in our analysis of the prostate cancer trial, and were nicely addressed by the discussants. In the future, promotion of DTRs in clinical studies will require close collaborations between statisticians and clinicians to overcome practical barriers and develop consensus guidelines for design and analysis.

Acknowledgments

We thank the editor, Hal Stern, for inviting this discussion. We also are grateful to the two sets of discussants both for providing many insights into the issues arising in our paper, and for pointing the way to new and important methodological research in dynamic treatment regimes (DTRs).

Additional References

- Bang H, Robins J. Doubly robust estimation in missing data and causal inference models. *Biometrics*. 2005; 61:692–972. [PubMed: 16135020]
- Bekele BN, Thall PF. Dose-finding based on multiple toxicities in a soft tissue sarcoma trial. *Journal of the American Statistical Association*. 2004; 99:26–35.
- Glasziou PP, Simes RJ, Gelber RD. Quality-adjusted survival analysis. *Stat in Medicine*. 1990; 9:1259–1276.
- Houede N, Thall PF, Nguyen H, Paoletti X, Kramar A. Utility-based optimization of combination therapy using ordinal toxicity and efficacy in phase I/II trials. *Biometrics*. 2010; 66:532–540. 2010. [PubMed: 19673865]
- Cao W, Tsiatis A, Davidian M. Improving efficiency and robustness of the doubly robust estimator for a population mean with incomplete data. *Biometrika*. 2009; 96:723–34. [PubMed: 20161511]
- Robins JM, Ritov Y. Toward a curse of dimensionality appropriate (CODA) asymptotic theory for semi-parametric models. *Statistics in Medicine*. 1997; 16:285–319. [PubMed: 9004398]
- Robins JM. Robust estimation in sequentially ignorable missing data and causal inference models. *Proceedings of the American Statistical Association Section on Bayesian Statistical Science*. 2000:6–10.
- Rotnitzky A, Lei Q, Sued M, Robins J. Improved double-robust estimation in missing data and causal inference models. *Biometrika*. 2012 To appear.
- Rotnitzky A, Robins JM. Semiparametric regression estimation in the presence of dependent censoring. *Biometrika*. 1995; 82:805–820.
- Tan Z. A distributional approach for causal inference using propensity scores. *J Am Statist Assoc*. 2006; 101:1619–37.
- Tan Z. Comment: Improved Local Efficiency and Double Robustness. *Interntl J of Biostatist*. 2008; 4:article 10.

- Tan Z. Bounded, efficient and doubly robust estimation with inverse weighting. *Biometrika*. 2010; 97:661–682.
- Thall PF, Wooten LH, Logothetis CJ, Millikan R, Tannir NM. Bayesian and frequentist two-stage treatment strategies based on sequential failure times subject to interval censoring. *Statistics in Medicine*. 2007; 26:4687–4702. [PubMed: 17427204]
- Thall PF, Cook JD. Dose-finding based on efficacy-toxicity trade-offs. *Biometrics*. 2004; 60:4–693.
- Zhao H, Tsiatis AA. A consistent estimator for the distribution of quality adjusted survival time. *Biometrika*. 1996; 84:339–348.