HHS Public Access

Author manuscript

Clin Trials. Author manuscript; available in PMC 2016 January 22.

Published in final edited form as:

Clin Trials. 2014 August ; 11(4): 418–425. doi:10.1177/1740774514538553.

University of Pennsylvania 6th Annual conference on statistical issues in clinical trials: Dynamic Treatment Regimes (morning session)

Keaven Anderson, **Marshall Joffe**, and **Michael R. Kosorok**

Keywords

Dynamic Treatment Regimes; Adaptive Trials; Statistical Issues in Clinical Trials; Sequential multiple assignment randomized trials (SMARTs)

Moderator: Andrea Troxel, ScD, University of Pennsylvania

Our first panelist is Keaven Anderson, who heads biostatistics for late-stage oncology drug development at Merck Research Laboratories, and among many other areas, has interest in adaptive trials with a focus on group sequential designs. Marshall Joffe is Professor of Biostatistics here at the University of Pennsylvania, and works in a number of areas involving causal inference, including different kinds of confounding and sensitivity of inference to temporal ordering. Finally we have Michael Kosorok, who is Professor and Chair of Biostatistics and Professor of Statistics and Operations Research at the University of North Carolina at Chapel Hill. Among his many research interests are statistical methods for clinical trials and personalized medicine. Each of our panelists will give us a brief reaction and discussion of the material we have heard thus far today.

Keaven Anderson, PhD, Merck

Thanks to the organizers for bringing this topic forward and the speakers for giving us excellent presentations. I will talk about where dynamic trials may be useful. One area that immediately comes up is oncology and understanding overall survival and the impact of drugs on overall survival. A dynamic regime design may be a nice way to gain some control

Conflict of Interest

None declared.

Participants

Andrea Troxel, ScD, University of Pennsylvania
Keaven Anderson, PhD, Merck
Marshall Joffe, MD, MPH, PhD, University of Pennsylvania
Michael R. Kosorok, PhD, University of North Carolina, Chapel Hill
Elizabeth Kumm, MA, MS, inVentiv Health Clinical
John Whyte, MD, PhD, Moss Rehabilitation Research Institute
Philip W. Lavori, PhD, Stanford School of Medicine
Maha Karnoub, PhD, Calgene
Estelle Russek-Cohen
Maha Karnoub, PhD, Calgene
Susan Ellenberg, PhD, University of Pennsylvania

of how patients are treated after the primary treatment in a trial, allowing a simpler causal inference assessment of the impact of the initial regimen.

One question that comes up is: how many stages do dynamic designs have to have to get such information? In the pharmaceutical industry, what we tend to think about is getting a drug approved and then everything else is relegated to a lower level of importance. Given that priority, can we generate interest in SMARTs and how we might invest in them? One question may be whether SMARTs could be done through large collaborations. I don't know if you would call the I-SPY trial a SMART, but it is a very complex collaboration of a lot of people and companies with different people investing in and leading it. This may provide an investment model for a SMART. Alternatively, is there a simpler model for a pharmaceutical company to consider? One example that has come up recently is that oncologists prefer sequential monotherapy treatments over combination therapies for some indications. Our usual design is to add an experimental drug on top of a standard control treatment. Having a dynamic design that allows comparison of sequential monotherapies versus a combination therapy seems like a potentially interesting paradigm. What I am trying to think about today is how can we take simple dynamic designs forward initially?

Marshall Joffe, MD, MPH, PhD, University of Pennsylvania

I will be making a few comments related broadly to the talks that were given during the previous session. First I would like to discuss what is new in the current focus on adaptive treatment designs and treatment strategies and I will also have some questions embedded here for the speakers and perhaps for others. First, to elaborate on some of what Dr. Lavori said in the first talk today, adaptive treatment strategies have been used previously in trials and in practice and in some ways even as part of strategies and trials. Consider two examples. One would be a study for a treatment of high blood pressure. It would often be the case that a protocol would specify changing the initial agent if the blood pressure is too high. That sort of a protocol is an adaptive treatment strategy as well. This can be compared with the standard treatment assuming that there is low non-compliance, but even if there isn't, using intent-to-treat sorts of analyses. Even though here there was an adaptive strategy, the adaptive part of the strategy wasn't typically a major focus of things. Similarly, in trials of many pharmacologic agents, the protocols themselves, which can be viewed as strategies, will allow for changes due to side effects or adverse reactions.

The earlier methods and trials, however, before the relatively new literature over the last decade or so were not really tailored for evaluating second line therapies. But you can use a standard design in analysis for evaluating adaptive strategies even when the adaptive part is part of what you are interested in. As an example, consider a trial with four arms, where there are two treatments, an initial treatment and a subsequent treatment. The second line treatment can be adaptive. In this case, if it is a treatment for blood pressure, version one could be to add a second treatment for blood pressure if the systolic blood pressure is greater than one target, let's say 140. A different strategy would add the second treatment if the systolic blood pressure is higher than a different target, let's say 160. Then you could use a standard design with four arms, A^1B^1, A^1B^2, etc., to compare those four different strategies, each one of which is an adaptive strategy. The focus here, in part, is on the adaptation, and

you can use a standard analysis using intent-to-treat rather than newfangled fancier analyses. What this is not is a way to optimize what to do. It allows you to evaluate four different options and essentially that is it. It doesn't let you choose the best, let's say in this case, the best blood pressure to target adding a second treatment. An issue common to all of the talks, is the comparison of multiple regimes or plans. Some of the speakers used the term "tailoring variables." If we take this down to simple clinical trials for simple treatment, rather than adaptive treatment, this might be known as a subgroup analysis. It is now extended and made more complicated by reference to multiple treatments and times and dynamic regimes. This sort of thing in the standard clinical trials literature often has an explanatory flavor, rather than a confirmatory flavor. In parts of the clinical trials community, especially the regulatory part, an exploratory approach often has a negative connotation. I come from more of an epidemiology background, where people do this sort of thing all the time and don't worry about it so much. One question, might be how these sorts of ideas might be incorporated into regulatory sorts of trials, or regulatory sorts of decisions under the circumstances where there are pre-specified comparisons. One idea, which was brought up in the first talk, was in the context of learning systems, which are outside the regulatory framework and perhaps that might be a more natural place to consider these ideas.

Issues of multiplicity were dealt with a little bit. The examples that were provided by the speakers and the simple example I gave deal with no more than one or two covariates, but in principle, there may be many covariates at each different point in time; how do you choose covariates and the number of covariates for the tailoring strategies? This was discussed a little by one of the speakers, in terms of eliminating something that doesn't seem to be terribly important. But if the trial is too small, by doing this you could end up wrongly removing a true tailoring variable from consideration. From a different literature there is a concern with overfitting. It is also well-known that if you fit and then evaluate a model using the same data from which your estimates were derived, they will be overly optimistic. In this context, how would you deal with the over-optimism of estimates of the predictive ability of the value of a particular treatment strategy?

Many of the methods presented are variants of methods that have already been developed for time-varying treatments and covariates, especially by Jamie Robins and co-authors, including the g-formula, inverse probability of treatment weighting and marginal structural models, which often go together, and then g-estimation, which was mentioned here, but people haven't discussed actually using it. That, for better or worse, seems to be the case not only in the adaptive trials literature, or adaptive regimes literature, but also in the causal inference literature, and I would be interested in hearing from people who do more of the adaptive treatment strategies and hear why that might be the case in this context. One of the new things in this literature is a focus on determining optimal regimes and strategies. As has been discussed, it has resulted in non-regularities and some solutions have been presented in this active area of research. For these sorts of trials you can randomize a strategy up front, or you can do multiple randomizations. The analysis strategies for these two types of designs in principle should be the same or similar, but there may be some practical differences between the two, and I would be interested in hearing what people have to say about that. It would also be worth having a little more discussion of approaches for selecting covariates for tailoring, both from observational studies, where we can apply any of the general methods

for looking for effect modification. Even in simple randomized trials we can analyze these in various ways: some might analyze akin to randomized trials or like observational studies, and some might take advantage of the initial randomization and use different sorts of assumptions. In this context, can one take individual preferences into account and maximize, for an individual, that person's own individualized utility functions, rather than come up with a broad recommendation? That would be presumably less the case if survival is the main outcome, but if other things are of concern one might want to do that. Finally, in this literature there is an emphasis on decision making and less on understanding of causal processes. Can some of those ideas of understanding processes be worked into this literature as well?

Michael R. Kosorok, PhD, University of North Carolina, Chapel Hill

I liked the gradual progression of the talks and the introduction to the whole concept of adaptive treatment strategies, dynamic treatment regimes, their practical importance, and then the three recognized challenges in estimating dynamic treatment regimes (DTRs). I am going to follow up with some of these and other approaches, alternatives to data analysis, in addition to Q-learning, that may be more efficient in some settings and also the challenges of inference in this very complex situation. Just the mere fact of its difficulty is important to note. Some approaches seem to be successful, but all of them are non-trivial for sure.

I want to talk first about some of the practical points that were raised. Adaptive treatment strategies are what doctors are trying to do and what patients want doctors to do. We want them to adapt based on what they learn about patients. Many of the young physicians that I interact with are quite eager to get SMART designs going because they allow them to test the strategies they are interested in knowing about. They are eager, even though these are very different from traditional approaches, and while young researchers are very anxious to use them, senior researchers are terrified of them. It really depends on the clinical research setting and the culture. I found this to be very interesting, because I work with people in the School of Nursing, versus people in the Cancer Center, versus people in cystic fibrosis and other areas and everybody has a different way of reacting to these new designs. I think we are still in a learning process of how to work with DTR protocols and get them through IRBs. I don’t have any magic answers here. I am just finding a lot of interesting experiences with it.

There are a lot of very important research opportunities for us. There are still many challenges to design DTR studies in efficient, practical ways. Obviously, there are also estimation challenges and inferential challenges. There are communication challenges. Somebody could spend their time just figuring out how to explain these ideas. Are they ready for prime-time? We have a lot of theory. We have a lot of simulations. We have some test cases in some disease areas, usually not so much in clinical areas like cancer, but in other areas, like addiction and depression treatment. Are there concerns about the fact that we have different challenges with inference? I think we are presented with an opportunity for a culture change.

SMART designs are complex. They are very different from traditional designs and if we do them right they may yield a “black box” treatment for which we have no understanding of underlying mechanisms. You give a treatment rule to the physician and they plug in the data and the rule will tell then what treatment to give. This approach can be shown in some situations to be the absolute best thing to do. You can already see the regulatory and other challenges to this approach. Many researchers prefer knowing a mechanism before they proceed with a new treatment, but unfortunately that would significantly delay the benefit of these treatments. We actually have examples throughout the history of medicine where we didn’t know the mechanisms initially. The early, very successful treatments of tuberculosis had no known mechanistic basis for efficacy. Later on we eventually understood the mechanism, but it took a while. Most current research most likely involves complicated mechanisms, and if we wait until we figure it out we will all be dead. Unfortunately, SMARTs can be hard to sell to key decision makers and funding groups, and also IRBs, but we should press forward with the assurance that these have potential to improve human health. We do have to be able to explain these things better, but I think it is very important to recognize that we don’t have to wait to implement something until everyone understands it. Otherwise, none of us would be using cell phones. It is very important that we somehow bridge that chasm where the experts in different areas of medical research are allowed to create technical innovations that people on the other sides of the chasm, or other key shareholders, don’t have to understand everything that goes on with it.

I wanted to just comment briefly on some of the points that were raised by the previous discussants. How many stages do you need for this technique to be useful? One is probably good, in the sense that some of the issues for simple or one-stage decision making have not been fully worked out and can be improved upon. For example, if you do a randomized trial and collect a lot of biomarkers, you can actually re-analyze the data, including dosing information, and go back and say, okay, if we tailor it, we see a significant benefit. You can come up with very simple examples where (and this is more hypothetical then real) a treatment for one group is good, but for another group it is toxic. If you do a traditional randomized study these effects will wash out and you won’t see anything, but if you use these adaptive methods, you can find the effects. Using a treatment in the right way for the right people is actually much better than just using standard of care. However, even with two stages in a dynamic design you can see dramatic improvements over more traditional methods because then you have an opportunity to adapt to the patient. Just one adaption makes a big difference.

Another point that was discussed is the role of p-values in these designs. We have to be willing to change the way we view evidence criteria. We definitely need the results to be reproducible and reliable, but it may be that focusing on confidence intervals may actually delay some of the results. For example, some of the black box approaches that are used to estimate effects are known to capture very complex structure in a very reliable way, but we have no idea how to do inference yet. One thing to keep in mind is that you are basically in discovery mode with these trials. You are not doing confirmation. Inference is important, but if we can identify likely treatment effects, and we know theoretically that these results are reproducible, then we can do a well-designed confirmatory trial with more standard inference. It may be hard to try to do these both at the same time.

Another interesting issue is the culture of subgroup analysis. In the setting of SMARTs we really want the patients to be different or else our adaptive strategies are not going to make any difference. We are essentially looking for subgroups who respond to treatment differently. This is a principled approach to subgroup searching. We are not just going on a fishing expedition. We are letting a procedure with established operational properties look for patterns. That may still be unsettling for some, but at least we know how the process behaves. Also, cross validation ends up being a very useful tool in these settings. It is something that we don't use as often in biostatistics as we should, but it is a very powerful tool that can help address some of the problems that can occur when you try to estimate something with the same data that you used to generate the estimates.

I want to comment about Q-learning related to g-estimation. One of the key benefits of Q-learning is that you can bring in a lot of machinery from a lot of different areas right into the problem. You can take everything from the machinery and use it, including support vector regression, or you can use randomized forests; all of these very powerful non-parametric or parametric methods, can be applied. There can be a lot of efficiency gains as well, not always, but there are some. There are some real advantages to Q-learning over g-estimation, although I think g-estimation could potentially be adapted in the same way, but it is easier to see how to do it with Q-learning.

A previous comment indicated that you could pre-randomize everybody in a dynamic treatment design. While this is correct, for many of the sequential treatment decisions it gets harder and harder to do complete pre-specification. If you have treatment choices depending on whether a person responded or didn't respond, but there exist certain conditions in which there are only certain treatments you are allowed to randomize to, it can be very complex to prescribe all such in advance. It is sometimes easier at the point of treatment decision to do the randomization. I also think that there may be an advantage for the people who are giving the treatments to not know the full sequence.

I also want to mention some recent work by Dr. Laber [3] on how to construct user-specified utilities. These could be used to decide whether you have a personal tolerance for one particular toxicity at the expense of survival. This is a very new and open area of research. These are just some of my thoughts on this very exciting area, where there are a lot of opportunities and challenges. I think the environment is similar culturally to when group sequential methods were first introduced. Although I think this new situation is a little more complicated than that, there are some interesting opportunities for us to be involved in a significant cultural shift.

Questions from the Floor

Elizabeth Kumm, MA, MS, inVentiv Health Clinical

I would like to focus in on what we mean by the word "optimize" because it sounds to me like what is being said is that through the course of the trial each person has received what was optimal for them and that's very different from what industry wants to argue when it goes before the Food and Drug Administration. They want to be able to say, 'My drug is

better than this drug.' That's a slightly different claim from, "Everybody optimized their treatment plan through the course of their disease."

Michael R. Kosorok, PhD, University of North Carolina, Chapel Hill

One way to think about this is to think about a regimen as a treatment. A regimen means a series of rules that at each point in time for a patient with these measurements, you give him/her this treatment. You have a complete rule for all of the patients. That's your treatment. The question is whether this is superior to another treatment or if you give a particular drug or standard of care to everybody. You still have, when all is said and done, a new treatment. It is just that it is not a single drug now. It is a set of rules to apply to the patients. One can actually imagine a phase 3 trial where you start it off with a SMART design and you add a bunch of fixed treatment alternatives. You could take the best of the fixed treatment alternatives and then you could compare them to the personalized tailored best. One would expect that this would be better than those, but you could at least test it, because it is not necessarily the best, and where there is one truly best treatment path for everybody, there is no reason to tailor, but if the tailoring creates a benefit then you want be able to use it.

Elizabeth Kumm, MA, MS, inVentiv Health Clinical

We are going to have to develop a better way to define what we mean by 'better' if better is one thing for one person and different for another.

Michael R. Kosorok, PhD, University of North Carolina, Chapel Hill

We may have to work harder at coming up with a definition of what it means for one treatment to be better than another. I think that one approach is to define "better" in terms of the expected population average of the clinical outcome for the treatment with everybody included. There may be better ways. That is a little bit different from clinical trials in which you have a defined group and you say this drug is better than the other drug in this defined group. You want to bring that up to maybe a more pragmatically selected group of people that are heterogeneous and say, in this population, if you apply this treatment rule, their outcome will be this expected value, whereas if you use other treatment rules it will be this other expected value.

Keaven Anderson, PhD, Merck

I think what's probably missing is how do you show that my new drug actually contributed benefit, as opposed to being part of an optimal treatment regimen. That's where we probably get stuck in the regulations for drug approvals.

John Whyte, MD, PhD, Moss Rehabilitation Research Institute

One of the issues that we face in neurologic recovery is that often there are not outcome measures that are well-suited to the entire spectrum of treatment changes. Are there approaches that would allow one to not use the same measurement yardstick at all decision points or does all of the underlying math assume that it is the same measurement tool that's being used throughout the adaptive design?

Michael R. Kosorok, PhD, University of North Carolina, Chapel Hill

The answer is no. Consider the following example. Suppose you have a single outcome at the end that may not be accessible at the decision times for treatment change, like overall survival or some other measurement of clinical status (where you may want to wait until six months after you are done with the trial). That's not a problem. The way that this works is that you actually do this estimation recursively backwards, by starting with the last estimation. At that last point you do have the outcome available and then you create a pseudo-outcome going forward and you are able to do the estimation even though you don’t have these outcomes at the earlier time points.

Philip W. Lavori, PhD, Stanford School of Medicine

I heard a slightly different question: whether the states that are being used to drive the treatments have to be measured on the same scale all along. The answer is no, which I think is also Dr. Kosorok’s answer. You have a synoptic outcome at the end, which is Y. Y is the summary of how the patient did over the whole trial period. It could be overall survival, or the average of some scores over that time, but it is the view that you only get at the end when you look back. Then the individual decisions along the way are made on the basis of what we might think of as myopic outcomes. I can see how things are going so far. I can see for example, my blood pressure at each time point, but I don’t know whether I have nephropathy. The states should be the most up to date things you know about at the time of a treatment change decision. What should be measured would depend completely on what's relevant to that treatment decision. Your scales could be completely different along different branches at different times, which I think is what Dr. Whyte asked.

Maha Karnoub, PhD, Celgene

What is the nature of collaborations with collaborators having different motivations? The industry motivation in dynamic treatment regimes is to prove that a drug is doing better in this context, and this drug is doing better in this context as part of this treatment regimen.

Philip W. Lavori, PhD, Stanford School of Medicine

The business model in the past has been that you develop and produce a drug that works on average for a very large number of people, including some number of people for whom it doesn’t work at all, and a number for whom it works wonders. If you sell it to a very large number of people, you recoup your costs and make a profit. There is a newer business model that says, no, we should do genetic testing and make treatment decisions based on a patient’s biomarker status along with other salient factors. Unfortunately, this is a really a non-intuitive approach. If you get a signature that's based on 100 variables that came out of a support vector machine with cross validation and so on and so forth, no one has intuition about that. You can't look at it and say, oh, yeah, I see. That's this pathway and that pathway. So black boxes seem to be the way of the future. But a better black box doesn’t just include what happened in some meiosis inside of my dad and my mom back before I was conceived and that led to my genotype, but also what happened along the way, like yesterday when I took the drug. That's another kind of personalized medicine based on what happened after treatment. Did I respond to that drug or not? Did my biomarkers go up or down? That's

coming. If you think the whole idea of doing genotype-based fractionation in the market is upsetting the apple cart for industry, wait until you see what happens with the biomarkers measured over time, such as human immune monitoring and seeing what your T-cells are like for every minute of the last 15 minutes. I think that if we can imagine where the world is going, it is going in the direction of personalized medicine, or precision medicine, and the DTR idea just adds the time dimension to it. In fact, in statistical terms it may turn out that if you measure enough good stuff now after you've given some treatments, the future is Markovian with respect to your genotype. In other words, the genetics may get caught up in some sense. Depending on where you see the future going this is either something that you have to adapt to in industry or it is something that you are going to try not to adapt to as long as possible.

Keaven Anderson, PhD, Merck

I will just say in defense of industry that I think we are greatly in tune with the fact that personalized medicine is really the key to the future. We are scratching our heads about biomarkers and we will continue to do so for a long time. I think this is another full employment area for statisticians for many years to come.

Philip W. Lavori, PhD, Stanford School of Medicine

Right now the current state is that when you utilize Q-learning and dynamic treatment regimes you already have candidate biomarkers and candidate drugs to feed into it and so you have already approved the drugs (e.g. you have already gone through the regulatory process). So the question remains as to whether you are using these new designs to validate a drug or to get it approved. Right now we spend a lot of effort to narrow down the drug choices and biomarkers ahead of time and then proceed with Q-learning or dynamic treatment regimes. But suppose you actually can design drugs that directly feed into this process. That might be a much more efficient approach.

Estelle Russek-Cohen, PhD, Food and Drug Administration

My personal view (not necessarily the view of the FDA) is that you don't always understand how every product works in terms of the underlying science. As Dr. Anderson says, you get out the door with an initial indication, but we realize that it is not the only way the product is going to be used later. I do have concerns about some of the exploratory investigations and whether they lead to grossly off-label use and what the consequences night be, but that could be judged on a case by case basis. I just want to make sure that people don’t think that we understand 100% of what a product does because it goes to market.

Maha Karnoub, PhD, Calgene

I want to remind people how even the simpler trials are logistically complicated. We got into the biomarker business and now we are adding even more complications with DTRs and the like. Philosophically, ethically, and medically when we present these approaches to solving problems we also have to think about the practicalities of doing them. That's maybe where the resistance from the senior investigators comes because maybe they are aware of the logistics a little more than the younger investigators.

Susan Ellenberg, PhD, University of Pennsylvania

It is complicated, if you want to show that your cancer drug improves overall survival and people change their treatment once it appears that your drug may no longer be working. Even if your drug seemed to delay progression longer than the comparator, at the end survival may not look very different. There are complications whichever way you do it. The SMART design at least has the advantage of setting up some kind of structure that will allow you to evaluate your drug's effect.

I did want to make one comment on the issue of timing of randomizations that Dr. Joffe mentioned. Suppose you develop a strategy for moving from one treatment to another when the first one stops working, and you randomized everyone up front to follow different paths. This creates an important problem. When I talk to investigators about clinical trials I always try and encourage them to randomize as close as possible to the time at which they are going to start treatment, because if you randomize them several weeks before they are going to start treatment some people are going to drop out and it creates havoc with your analyses. It seems to me if you have multiple stages and you try to randomize up front to a particular strategy you are going to lose a lot of people along the way. It won't be as problematic for inference if you randomize at each stage because you will have separate randomized cohorts that can be analyzed in a more straightforward manner.

Michael R. Kosorok, PhD, University of North Carolina, Chapel Hill

Suppose you get to the second stage and you qualify for a randomization, but 30 percent of the people choose not to be randomized. I wonder if one of you would like to comment on the practical issues surrounding that.

Philip W. Lavori, PhD, Stanford School of Medicine

I think patients that choose not to be randomized after the first stage are a huge problem. Dr. Ellenberg is right about trying to keep the randomizations as close to the decision point of the treatments as possible because it helps with that, although it is not perfect. The missing data problem and failure to randomize is I guess bigger in these sequential randomization trials and certainly impacts the analyses and inference. I don’t think there is a magic answer. The trick is to figure out how to do the consenting and the randomizations in a way that makes it as easy as possible for people to just slide along. As usual the principle is you don’t want to try to randomize people to things that they are going to feel very differently about. It is very hard to randomize people to surgery versus medicine, for example. There may be some ways of ascertaining up front what someone's equipoise stratum is. There is another point. There are two schools of thought in SMART design, not opposed but differing in emphasis. One school of thought, which I think of as being from Dr. Susan Murphy and her colleagues and students, says it is a learning exercise and you are going to develop estimates of optimal strategy. That's a terrific way of thinking. There is another, more stick-in-the-mud sort of thinking, which I identify myself with, where you already have a pretty good idea of what the candidate strategies are-- I am thinking of something like the SMART study of rituximab to address early and late treatment failures in older patients with diffuse large B-cell lymphoma, using either R-CHOP versus CHOP for induction and then R versus nothing for maintenance in those who respond [4]. That trial had a very simple, well-identified

approach that a drug company might use because it fits into what they are trying to accomplish. In that situation, it seems to me that the SMART design just gives you a way of efficiently estimating the things you need to estimate, which is the comparison of these four adaptive treatment strategies. That's the confirmatory strategy. I think the FDA is content with these approaches in the same way that they would be content with an upfront randomization. There are efficiency gains, in general, if you don't do the upfront randomization. The study in some sense is more self-designed if you do the randomization, as Dr, Ellenberg describes. There are confirmatory SMARTs, after which you can take the gloves off and do lots of Q-learning if you want, and you should. Then there is the possibility of a developmental study, and Q-learning, or something like it, as the central focus. I think we need both approaches.

Michael R. Kosorok, PhD, University of North Carolina, Chapel Hill

This raises one of the strategies for protocols. What seems to be helpful is to actually couch the more personalized approach as a later aim in a clinical trial where you have randomization of these paths and your first aim is to find the best treatment path. Then you use the same data to go for the next aim and from that you come up with your treatment regime, for which you may actually want to design a later, confirmatory trial as well. One strategy is to couch these both together so everybody is kind of happy, although it tends to make everybody mad instead. In any case, the strategy is a good one. Some of the resistance that I have observed relates to traditions of data analysis in some groups that are extremely strong. I am thinking of nursing researchers I have worked with.

Some people in nursing have very traditional views of regression models. The alternative approach to the analysis we proposed just confused them and there are others who really want to see a coefficient, a parameter that you can estimate, both of which you would have in traditional trials. Both of those groups don't like something they don't understand statistically. It is not just a logistics issue. It is also I think important to be able to communicate and change the culture about what roles statistics and biostatistics play in this process.

Acknowledgements

The Center for Clinical Epidemiology and Biostatistics at the University of Pennsylvania provided scientific and logistical support and the American Statistical Association and the Society for Clinical Trials provided in-kind support for the conference.

Funding

Funding for the conference was made possible (in part) by award CA132565-06 from the National Cancer Institute. The views expressed in written conference materials or publications and by speakers and moderators do not necessarily reflect the official policies of the Department of Health and Human Services; nor does mention by trade names, commercial practices, or organizations imply endorsement by the U.S. Government.

Glossary

A Dynamic Treatment Regime (DTR)

individualizes treatment over time via decision rules that specify whether, how, or when to alter the intensity, type, or delivery of treatment at critical clinical decision points in individual patients [1].

A DTR is said to be optimal
if it yields the best clinical outcome on average when applied to the entire patient population [2]

Q-learning
uses approximate dynamic programming *to estimate an optimal DTR* with the requisite expectations being approximated with regression models. (from [2] slightly modified]

Sequential multiple assignment randomized trials (SMARTs)
or equivalently, sequentially randomized trials, have been developed explicitly for the purpose of constructing proposals for high-quality DTRs [1].

References

1. Almirall D, Lizotte D, Murphy SA. SMART Design issues and the consideration of opposing outcomes. Discussion of “Evaluation of viable dynamic treatment regimes in a sequentially randomized trial of advanced prostate cancer” by Wang, Rotnitzky, Lin, Millikan, and Thall. J Am Stat Assoc. 2012; 107:509–512. PMCID: PMC3607391. [PubMed: 23543940]
2. Chakraborty B, Laber B, Zhao Y. Inference for optimal dynamic treatment regimes using an adaptive m-out-of-n bootstrap scheme. Biometrics. 2013 Sep.69(3) PMCID: PMC3864701.
3. Laber EB, Lizotte DJ, Ferguson B. Set-valued dynamic treatment regimes for competing outcomes. Biometrics. 2014; 70:53–61. PMCID:PMC3954452. [PubMed: 24400912]
4. Habermann TM, Weller EA, Morrison VA, et al. Rituximab-CHOP versus CHOP alone or with maintenance rituximab in older patients with diffuse large b-cell lymphoma. J Clin Oncol. 2006; 24:3121–3127. PMID:16754935. [PubMed: 16754935]