Cytel

Topics in Bayesian Statistical Methods

DECEMBER 2022

Table Of Contents

Bayesian topics are frequently explored in Cytel's Perspectives on Enquiry and Evidence. What follows are a set of articles exploring various topics related to Bayesian statistical methods.

Sign up for our weekly newsletter to stay up to date on the latest thought leadership in data science.

Sign up for our weekly newsletter

Bayesian Methods: Paving the Path to Clinical Development Transformation

By Mansha Sachdev

Bayesian methods have been playing a key role in transforming clinical research in therapeutic areas such as oncology and rare diseases, and in addressing clinical development challenges for COVID-19 drugs, devices, and biologics. From early-phase trials to late-phase development, utilizing Bayesian tools can expedite and/or de-risk trials, even when used to augment a Frequentist framework. Yet access to such designs has been limited by the need for powerful computational modeling and deep statistical expertise.

Bayesian approaches provide a variety of new opportunities for efficient and flexible clinical trials. The Bayesian approach allows every new piece of data to serve as evidence to update a hypothesis. Bayes' rule consists of a "prior" either based on evidence already collected, or scientific findings in the case of early-phase trials. A rule then explains how to update these priors in order to make sense of newly collected evidence. A "posterior" is then the result of the prior being updated in light of this new evidence. As the trials evolve with new in-trial insights, these Bayesian methods enable statisticians and sponsors to create flexible trial designs and accelerate learning.

Frequentist designs can often require higher sample sizes than Bayesian methods and are considered by some to be less flexible and less intuitive. However, the decision to use Frequentist or Bayesian designs is a matter of context. After considering all the parameters, it is on the statistician's expertise to choose the method that is best suited to the objectives of the trial.

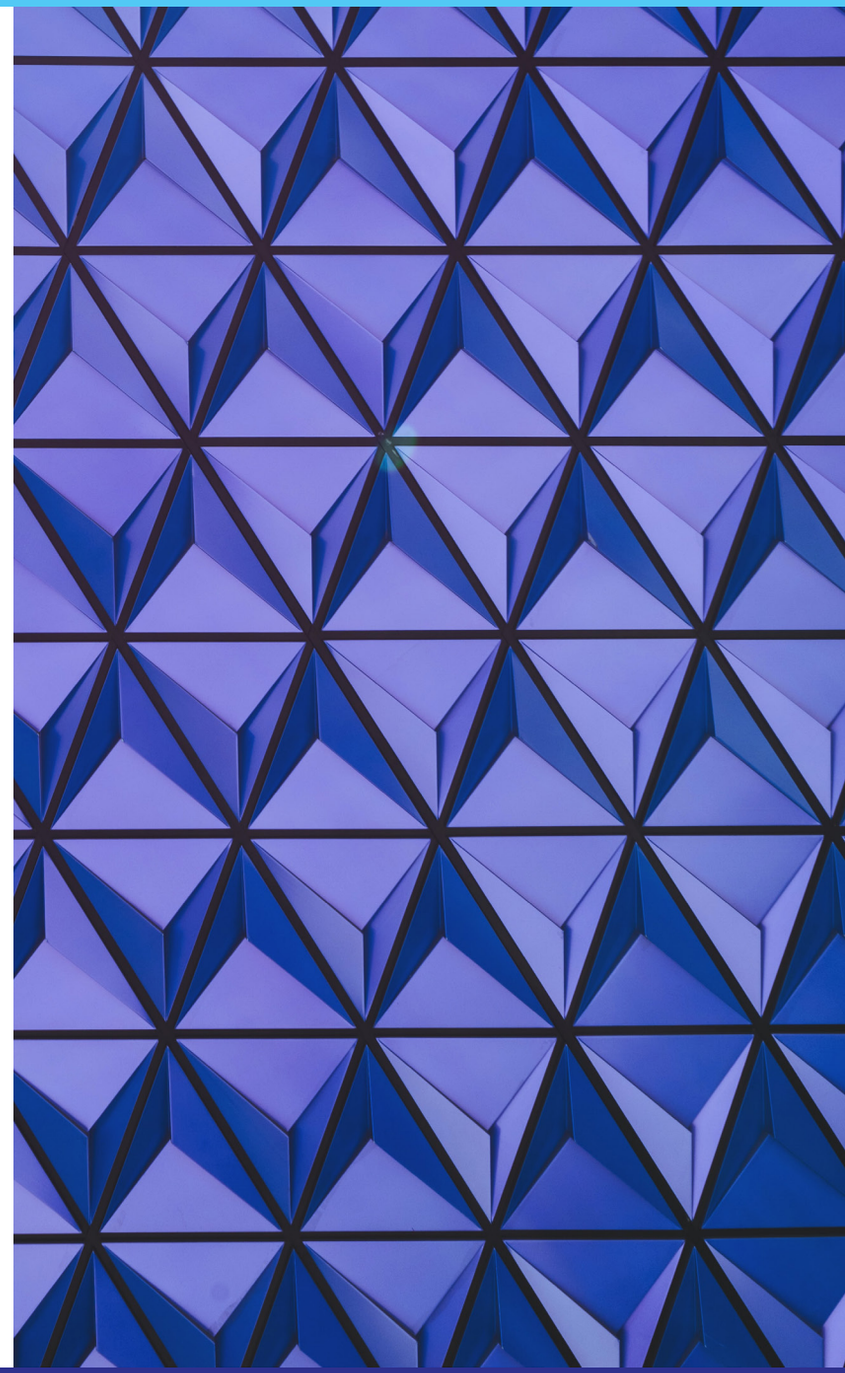

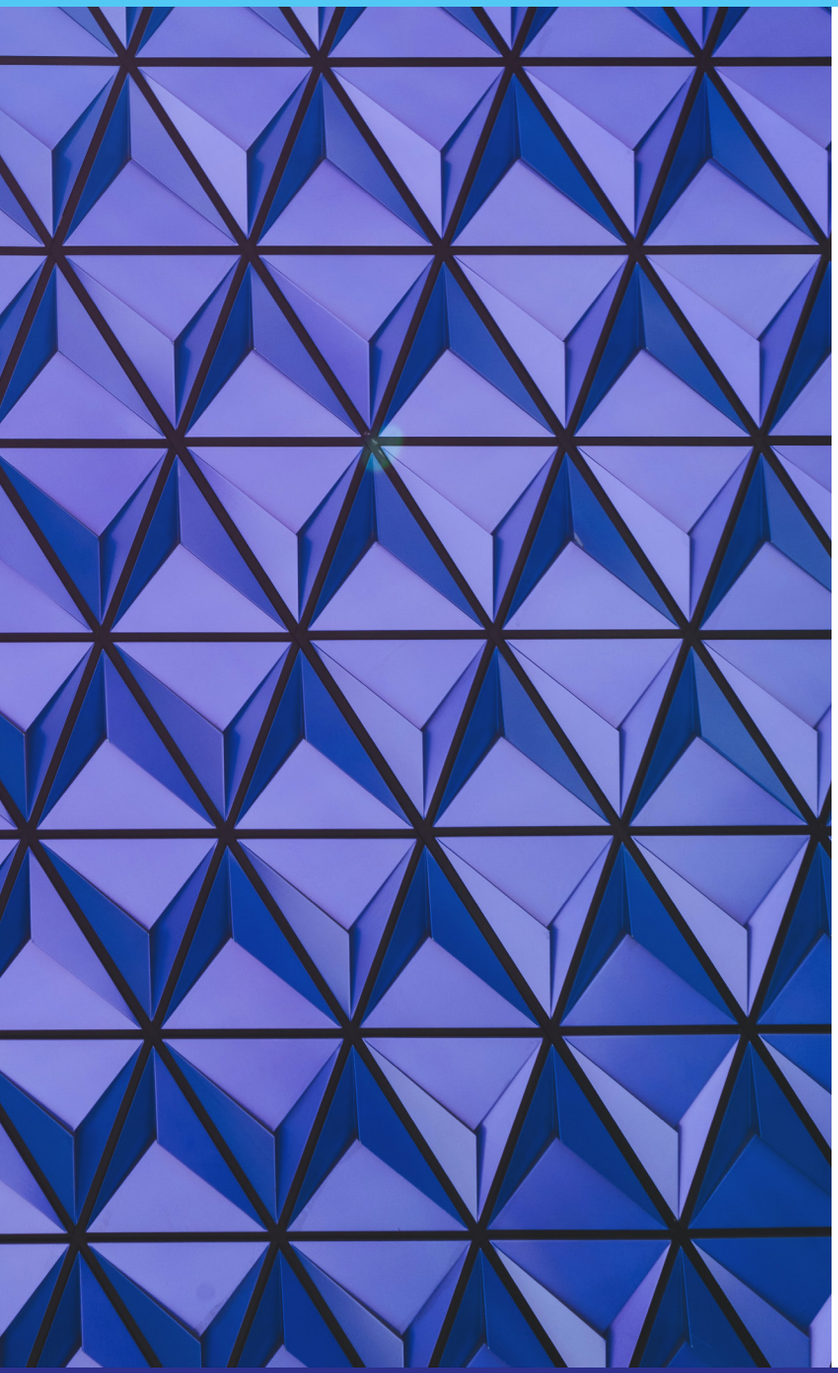

When taking a closer look at the context, scientists at Cytel have often advocated for mixed or hybrid methods to be used for a single submission. Not every sponsor realizes (and it is important to note) that Bayesian methods can also be used to strengthen insights of trials that are Frequentist. Suppose a Frequentist trial faces limited resources or needs a little more power to convince regulators of its statistical rigor. Bayesian methods can be used to borrow from historical datasets, incorporating already existing data into a new clinical trial to augment new findings.

Bayesian designs are frequently used in early-phase trials due to their flexibility and efficiency. As very little is known about the response of a drug or its toxicity, patients and sponsors benefit from a trial design that enables frequent interim looks. Bayesian methods provide the opportunities for adjusting doses and stopping for futility when required. They offer an intuitive approach to clinical development, maximizing the use of available information at each interim analysis. For fast recruiting trials, Bayesian designs offer high confidence in early futility, efficacy, and sample-size decision-making, basing the decision on the consistency of the results from two or more early interim analyses.

Bayesian methods can also be integrated with Frequentist clinical trial designs to obtain clearer benefit-risk profiles for a number of new therapies. Ethically, every ounce of data collected ought to be used for these calculations, and Bayesian methods allow this to occur.

To learn more about this topic, see “Bayesian Methods: Transforming the Future of Clinical Research,” published in the Journal for Clinical Studies by Cytel Senior Research Principal Ofir Harari, VP of Customer Success Pantelis Vlachos, and CSO Yannis Jemiai.

Read the original post.

On Frequentist and Bayesian Sequential Clinical Trial Designs

By Heather Struntz

In clinical trials, patient enrollment is often staggered, with data collected sequentially. When designing a clinical trial, it is usually advantageous then to plan for interim analyses, which take a look at the accumulating data and present the potential for modifying the trial. Appropriately placed interim analyses are crucial because once a clinical trial has achieved enough data to determine its success or futility, decision-makers can plan for an early stop, saving resources and avoiding unnecessary enrollment. Here we take a brief look at the fundamental and philosophical differences between two approaches to determining this critical stopping point in such adaptive clinical trials: frequentist and Bayesian sequential clinical trial designs.

Frequentist clinical trial designs

Frequentist probability, generally, understands the probability of an event as the limit of how frequently it might occur across infinite repetitions of an experiment. That is, for any event, it will either occur or it will not, and the frequency of it occurring, as observed across repetitions of an experiment, is a measure of that event's overall likelihood.

In clinical trials, frequentist designs employ repeated significance testing or conditional power to make early stopping decisions. Since parameters are viewed as fixed, a hypothesis regarding a particular parameter is either true or false; one cannot assign a probability to a hypothesis but instead can consider the Type I and Type II error rates. The Type I error rate refers to the chance of falsely rejecting the null hypothesis, given that it is true. Thus, frequentist designs need to be adjusted for the planning of interim analyses to achieve Type I error control, enabling interim analyses to be performed without error inflation.

Bayesian sequential clinical trial designs

Under the Bayesian paradigm, probability is interpreted not as a frequency, but as a degree of belief, in which the probability of a hypothesis is updated as more evidence becomes available. That is, with parameters viewed as random variables, to evaluate the probability of a hypothesis, a prior probability is specified, then later updated to a posterior probability with new data.

In clinical trials, Bayesian designs use these posterior or posterior predictive probabilities for decision-making. One can attach a probability to hypotheses, and thus the frequentist error rates are not necessary here to reject or accept a hypothesis—they are not central to Bayesian inference — though they may be calculated similarly even under a Bayesian framework. Early stopping rules are thus typically based on the posterior probability of an unknown parameter being greater than some threshold.

For an in-depth look into these design approaches, in particular in regards to early stopping rules for efficacy, as well as the numerical studies used to assess them, see Cytel executive advisor and University of Chicago Professor Yuan Ji's recent paper, "On Bayesian Sequential Clinical Trial Designs," arXiv:2112.09644.

Read the original post.

Use of a Bayesian Approach in Basket Trial Design

By Pantelis Vlachos

Advancements in biomarkers and momentum in precision medicine has paved the foundation for complex studies like basket trials. Basket trials are a type of master protocol, in which a targeted therapy is evaluated for multiple diseases that share common molecular alterations or risk factors that may help predict whether the patients will respond to the given therapy. Bayesian methods are particularly useful for these complex trial designs, as they enable greater flexibility and better ability to respond to the needs of the master protocol designs. Phase 2 Bayesian designs using hierarchical models, allow basket trials to efficiently assess the efficacy of a treatment in multiple disease indications.[1]

General Design Concept for a Basket Trial

A Basket trial is essentially a multi-arm Phase 2 or Phase 3 study investigating a treatment for multiple diseases or sub-diseases. They are usually conducted without randomized control. Normally, each arm in a basket trial is compared with a historical control. Patients enrolled in a basket trial are often composed of a heterogeneous group across multiple indications, such as different cancer types. Therefore, it is difficult to evaluate time-to-event endpoints, such as progression-free survival (PFS) or overall survival (OS). In basket trials, primary endpoints are often response rate (e.g., objective response rate (ORR) or pathological complete response (pCR)), which are less sensitive to the effects of population heterogeneity.

Basket trials often intend to use pooled population for primary analysis to gain broader indications across tumor types. However, clinical data to support pooling may be limited,

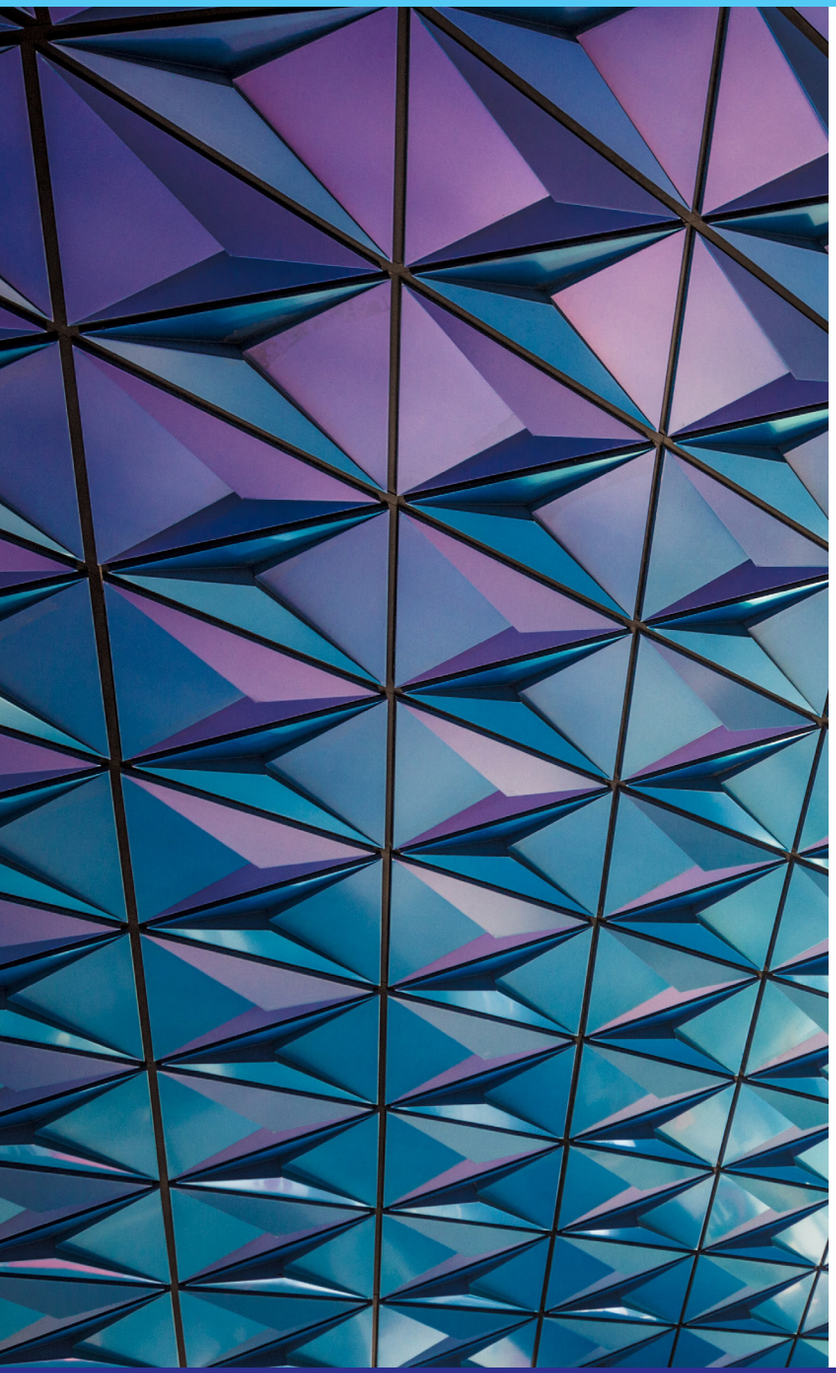

and treatment effect may differ between tumor types. For example, Vemurafenib works in melanoma with BRAF V600E mutation but not in colorectal cancer with the same mutation. When the homogeneity assumption is not valid, a separate stand-alone analysis for each arm is a simple alternative. However, conducting an independent evaluation in each arm is time- and resource-consuming. Also, the clinical trial sample size may be inflated under independent arms when compared to designs that borrow information.

Use of a Bayesian Approach

Recent publications have proposed adaptive designs that borrow information via model-based inference. Using the observed data, these methods borrow information by prior distributions that shrink the arm-specific estimates to a centered value. This is typically achieved through hierarchical models, where the shrinkage parameter, controlling the strength of information borrowing across different subgroups, is treated as an unknown parameter using a noninformative prior distribution.

In East Bayes, we implement a module of Basket Trial Designs and use simulation-based power calculation to evaluate four Bayesian approaches, including the Bayesian hierarchical model (BBHM) proposed by Berry et al. (2013), the calibrated Bayesian hierarchical model (CBHM) by Chu and Yuan (2018a), the exchangeability-nonexchangeability (EXNEX) method in Neuenschwander et al. (2016) and a novel multiple cohort expansion (MUCE) method in Lyu et al. (2020). Users may choose desirable designs based on the software provided in this module. Below are some benefits of each of these approaches:

1. **Bayesian Hierarchical Model (BBHM)**: The Bayesian hierarchical design is more likely to correctly conclude efficacy or futility when compared to Simon's Optimal Two-Stage

design, in many scenarios. It is a strong design for addressing possibly differential effects in different patient groups. As stated above, the shrinkage parameter, which controls the strength of information borrowing, is treated as an unknown parameter following a noninformative prior. The data is allowed to determine how much information should be borrowed across tumor subgroups.

2. **Calibrated Bayesian Hierarchical Model (CBHM)**: Chu and Yuan (2018a) proposed a calibrated Bayesian hierarchical model as an extension of BBHM. CBHM provides a practical approach to design basket trials with more flexibility and better controlled Type 1 error rates than the Bayesian hierarchical model. By linking the shrinkage parameter with a measure of homogeneity among subgroups through an appropriately calibrated link function, the CBHM allows information borrowing when the treatment effect is homogeneous across subgroups and yields a much better controlled Type 1 error rate than the BHM when the treatment effect is heterogeneous across subgroups.[2]

3. **Exchangeability-Nonexchangeability (EXNEX) Method**: The EXNEX approach allows each arm-specific parameter to be exchangeable with other similar arm parameters or nonexchangeable with any of them. This is achieved through a mixture model with three components: first, corresponding to exchangeable, efficacious groups; second, corresponding to non-efficacious exchangeable groups; and third, corresponding to nonexchangeable groups.

4. **Multiple Cohort Expansion (MUCE) Method**: MUCE design was originally proposed for trials with multiple arms, including basket trials. Built on Bayesian hierarchical models with multiplicity control, MUCE adaptively borrows information across patient groups from different indications treated with different doses. A hierarchical model accounts for the fact that when aggregating data across patient groups, some treatment arms might have more significant differences than others, and this might require statisticians to

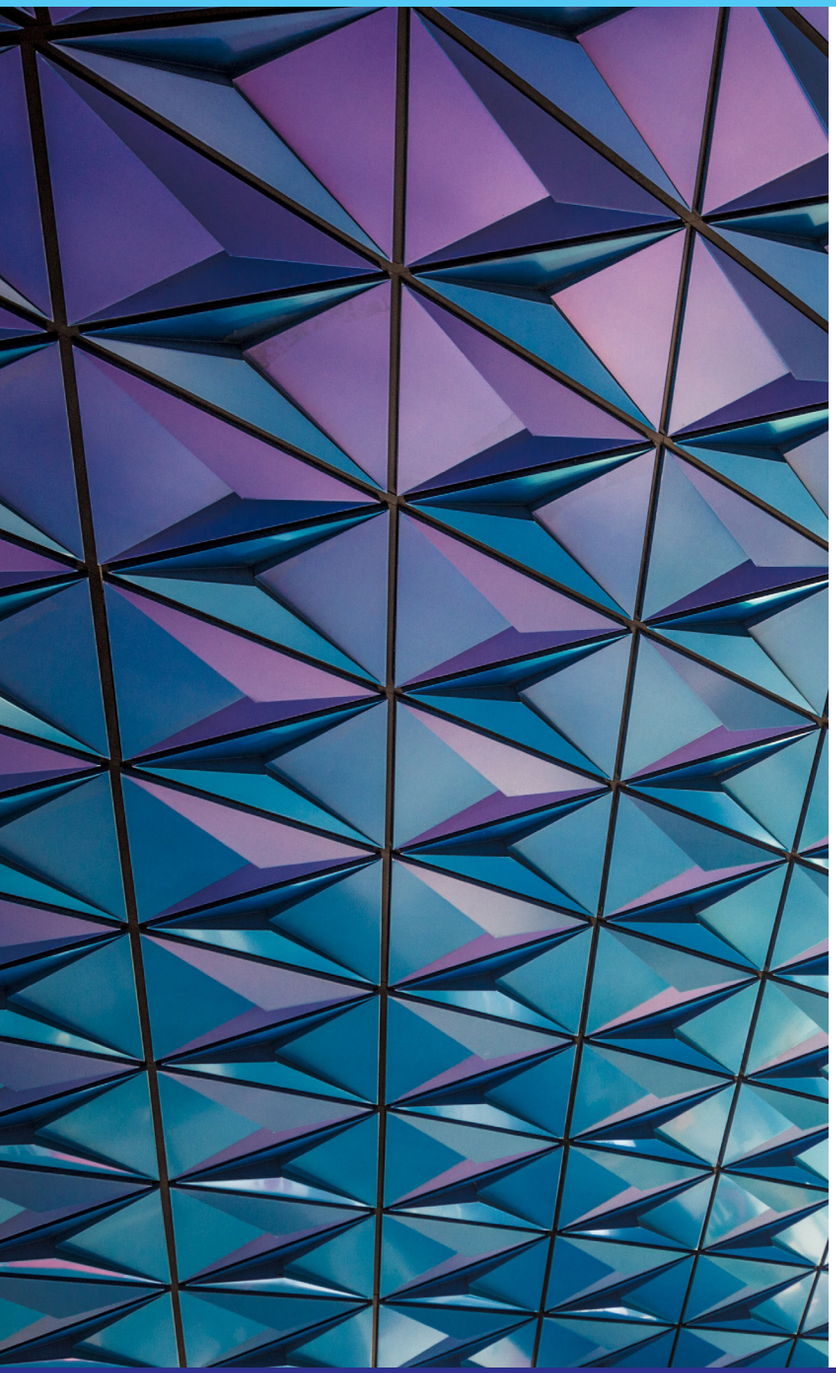

make adjustments by weighting to achieve unbiased measurement. This enables MUCE to control Type 1 error while increasing power and reducing sample size. These efficient designs can be applied in any clinical trials with two or more arms. For an expansion cohort trial in the United States, the MUCE design showed saving in sample size of up to 16.67% compared to Simon's 2-stage design.

Endnotes

1 Berry S, Broglio K, Groshen S, Berry D. Bayesian hierarchical modeling of patient subpopulations: Efficient designs of Phase II oncology clinical trials. Clinical Trials: Journal of the Society for Clinical Trials. 2013;10(5):720-734. doi:10.1177/1740774513497539.

2 Chu, Y. and Yuan, Y. (2018a). A Bayesian basket trial design using a calibrated Bayesian hierarchical model. Clinical Trials, 15(2):149–158.

Read more from the Informative Bayesian series.

Read the original post.

Introduction to Evidence Synthesis and Bayesian Dynamic Borrowing

By Pantelis Vlachos

In the last few years, there has been a growing interest in historical borrowing or augmented trials. There is an increasing level of comfort in using these methodologies even in confirmatory trials settings. The key challenge in borrowing external information is the selection of appropriate historical studies or external data sources. There are benefits to historical borrowing but also potential risks (for example, Type 1 error and power can be impacted by the drift).

However, despite the risks, several projects submitted to the FDA's Complex Innovative Designs (CID) initiative aim at using historical controls in Phase III studies. Many data-sharing initiatives such as TransCelerate, Project Datasphere, and others are all working toward making clinical trial data available for repurposing and reuse across the industry. There are also several working groups such as, the European EFSPI/PSI Historical Data Special Interest Group and DIA Bayesian Working Group who are interested in this area. Here we aim to introduce the concepts of evidence synthesis and Bayesian dynamic borrowing.

Evidence Synthesis

Partial extrapolation, multi-regional clinical trials and bridging studies are all forms of evidence synthesis. In this case, decision-making for the target group (for example, pediatrics, ethnic subgroup, or region) is based on the totality of evidence about overall treatment effect in the source population, combined with the consistency of treatment effects across regions/subpopulations and knowledge about intrinsic and extrinsic factors likely to impact treatment effects in different regions/subpopulations. Formal synthesis of

data from the source population and the target population using statistical modeling can help inform regulatory decision-making about regional treatment effects.

There are many ways to do evidence synthesis, but Meta-analysis is the most widely used method when there are several subgroups/regions of interest.

Bayesian Inference

Bayesian Inference is a form of evidence synthesis and Bayesian approaches provide a posterior probability distribution for some parameter (e.g., treatment effect), derived from the observed data and a prior probability distribution for the parameter. The posterior distribution is then used as the basis for statistical inference. For one of our clients, Cytel implemented a Bayesian framework of analysis, which allowed the data that had been collected during the clinical trials for adults to serve as “informed priors” for the pediatric trial. In Bayesian methodology, empirical evidence already available at the start of a trial is taken as the “prior” and then methodically updated throughout the course of the trial. Since Bayesian methods allow the use of data collected before a clinical trial commences, trial time diminishes.

Bayesian Dynamic Borrowing

They key thing to consider for the standard Bayesian approach is that if we use the source study data as a prior distribution then the posterior distribution for treatment effect is based on pooling the source and target study data. The source and target study data may not be entirely similar due to sampling variability or due to differences in inclusion/exclusion criteria between source and new studies, or health care may introduce differences between source and new outcomes. There can also be different populations/baseline risk factors

(such as ethnicity and age) involved in source and new studies.

The idea of dynamic borrowing is to account for the inconsistency between source data and target study population by learning how much information to borrow. The larger the drift, the less we borrow. The smaller the drift, the more we borrow.

Read more from the Informative Bayesian series.

Read the original post.

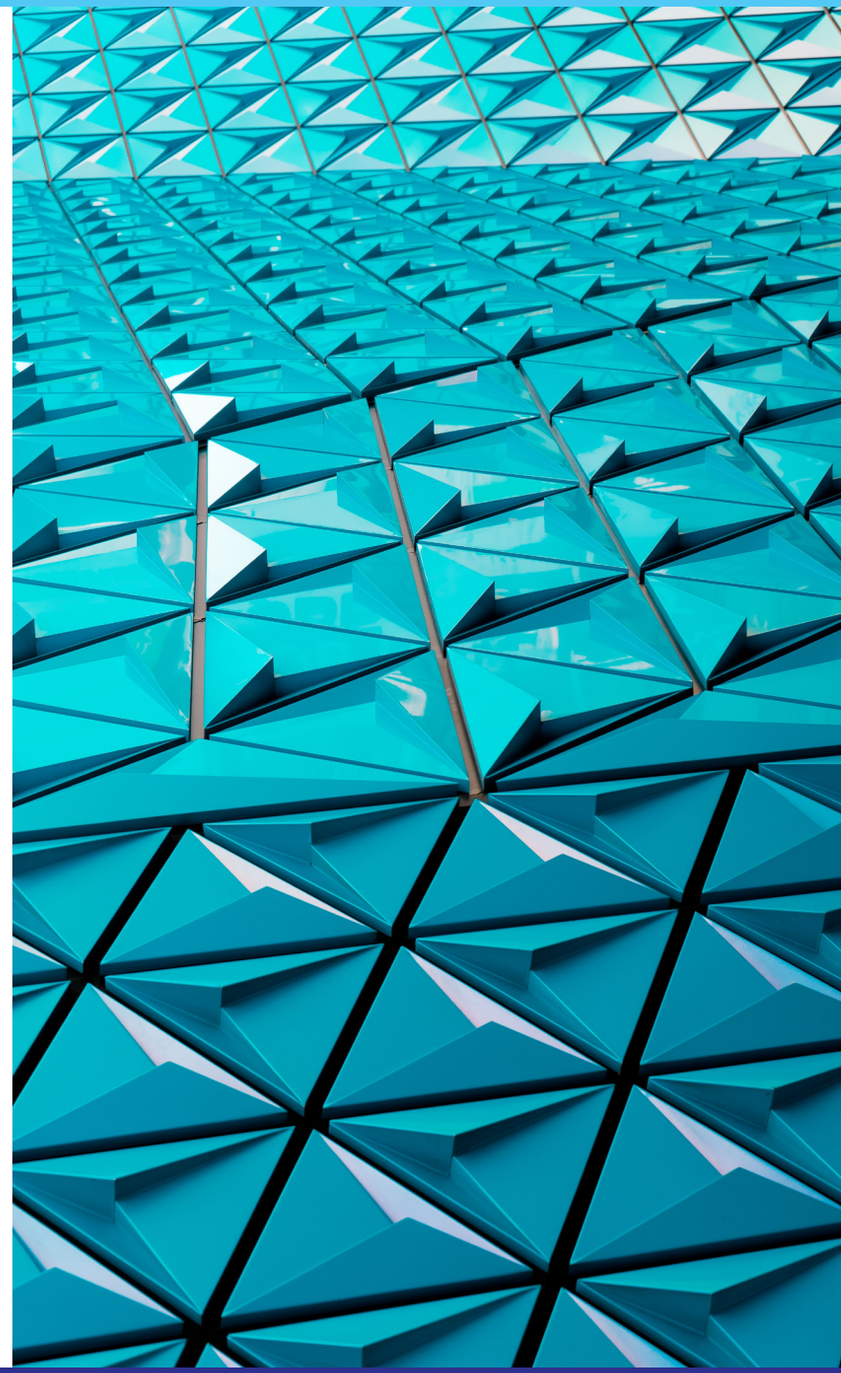

Bayesian Adaptive Clinical Trial Designs: INLA vs. MCMC

By Krishna Padmanabhan

How are Bayesian design models computed and fit analytically?

Typically, in Bayesian models, the objective of interest is to compute the posterior distribution or a predictive distribution. The standard method statisticians use to estimate these is called MCMC, which stands for Markov Chain Monte Carlo. MCMC methods are a class of algorithms for sampling from a probability distribution by constructing a Markov chain that eventually converges to the desired distribution at equilibrium. MCMC techniques have gained in popularity and adoption over the last few decades and have vastly influenced the uptake of Bayesian methods.

What are some limitations of MCMC?

While MCMC methods are reliable and can compute the posterior distributions associated with any likelihood function, they may be slow to converge and can take a long time to execute. Depending upon the complexity of the hierarchical model under investigation (the dimension of the unknown parameters), this may result in a sub-optimal exploration of the design due to computational demands and limitations.

What is INLA?

The integrated nested Laplace approximation (INLA) is a method for approximate Bayesian inference. It can be an attractive alternative to MCMC methods due to its speed and ease of use. Unlike MCMC, which relies on the convergence of a Markov chain to the desired posterior distribution, INLA uses a Laplacian approximation to estimate the individual

posterior marginals of the model parameters.

How does INLA work?

INLA uses a Laplacian approximation to estimate the probability distribution function of an unknown parameter by approximating it to a Gaussian distribution. The nesting occurs due to a posterior computation of the parameters conditional on the “data + hyperparameters” as a first step to computing the posterior distribution of the hyperparameters themselves. These two estimates are then used in a numerical integration calculation to compute the desired posterior marginal distribution. See the image below, for example, which shows the true posterior Gamma distribution in black and the INLA approximation to it in red. This example was based on a simple conjugate Poisson-Gamma example. The closer the posterior is to a normal-like curve, the more accurate the INLA approximation will be.

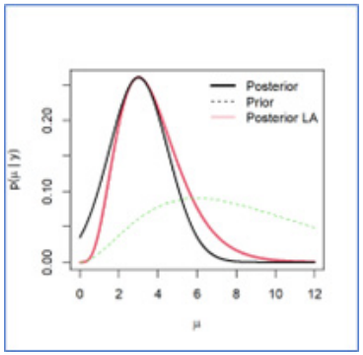

What are the limitations associated with INLA?

INLA only works for Latent Gaussian Models, whose parameters form a Gaussian Markov Random Field. The former is typically achieved by setting normal priors to (some transformation of) the unknown parameters. INLA, although seemingly limited to a certain class of models, works for:

1. Linear Models
2. Generalized Linear Models
3. Linear Mixed Models
4. Generalized Linear Mixed Models
5. Generalized Additive Models
6. Time to Event (Survival) Models

Thankfully, these models cover a large majority of Phase 2 and 3 clinical trials. It can also work for Phase 1 studies such as the BLRM, which utilizes a Bayesian Logistic Regression model.

Typically, how much faster is INLA than MCMC?

In Cytel's experience, we have noticed speed increases of >100x with INLA vs. MCMC in many instances, with virtually the exact same parameter estimates and confidence intervals. The graphic shows a comparison of compute times for different Bayesian clinical trial designs on a standard PC (Intel Core i7, 16GB RAM).

Comparison of Compute Times: JAGS vs. INLA			
Trial Primary Endpoint Type	MCMC (sec.)	INLA (sec.)	**MCMC:** 50K iterations, 3 chains **INLA:** Standard INLA Simplified Laplace Approximation
Survival (Oncology)	187	1.1	
Binary (Infectious Disease)	238	0.9	
Repeated Measures (Nephrology)	153	2.7	
Continuous (Rare Disease Biomarker)	36.2	1.0	
Survival (CV) (N=3000+)	>49K (13.7 hours)	27.35	
Repeated Measures (Lipids) (N=7000+)	>250K (~3 days)	396.2	

How does this help you design a better Bayesian clinical trial?

Although the increase in computational speeds is impressive, the real advantage is in a deeper exploration of design parameters, thus optimizing your Bayesian design further. Frequently, an MCMC explores limited design parameters due to computational demands associated with the process. Using INLA, a more through exploration of the design parameters can take place in order to achieve a greater confidence in the resulting trial design.

Read the original post.

The Uses of Bayesian Methods in Late-Phase Clinical Trial Strategy

By Esha Senchaudhuri

A number of late-phase clinical trial sponsors remain hesitant to employ Bayesian approaches in confirmatory settings, for fear that such statistical approaches generate obstacles for regulatory acceptance.

Bayesian methods, with their ability to facilitate flexibility and learning, are often associated with early-phase clinical trials. Their benefits for dose-finding and unplanned stopping in Phase I and II trials are clearly documented. Recently, more Bayesian elements have appeared in some confirmatory trials, particularly those requiring historical borrowing. A 2018 FDA Guidance on the uses of Bayesian methods in medical device trials appeared to normalize the idea that Bayesian methods can and should be used in some confirmatory settings. Pediatric trials, which modify existing adult therapies for younger populations, also benefit from Bayesian approaches.

Without doubt, the idea that Bayesian approaches provide strategic insight for clinical trials is not new. Neyman and Pearson's 1933 paper on hypothesis testing using Bayesian methods has long been the subject of debate, particularly regarding applications within clinical research. Steven Goodman argued in 1999 that several interpretations of statistical results were conflated by the use of p-values, and important distinctions about evidentiary meanings were more easily captured within Bayesian paradigms. Over two decades later, the debate remains unresolved.

Still, a critical element of late-phase regulatory submission requires sponsors to demonstrate Type 1 error control. In order to do this within a Bayesian clinical trial requires the introduction of Frequentist techniques in a Bayesian paradigm. Expert statisticians well-

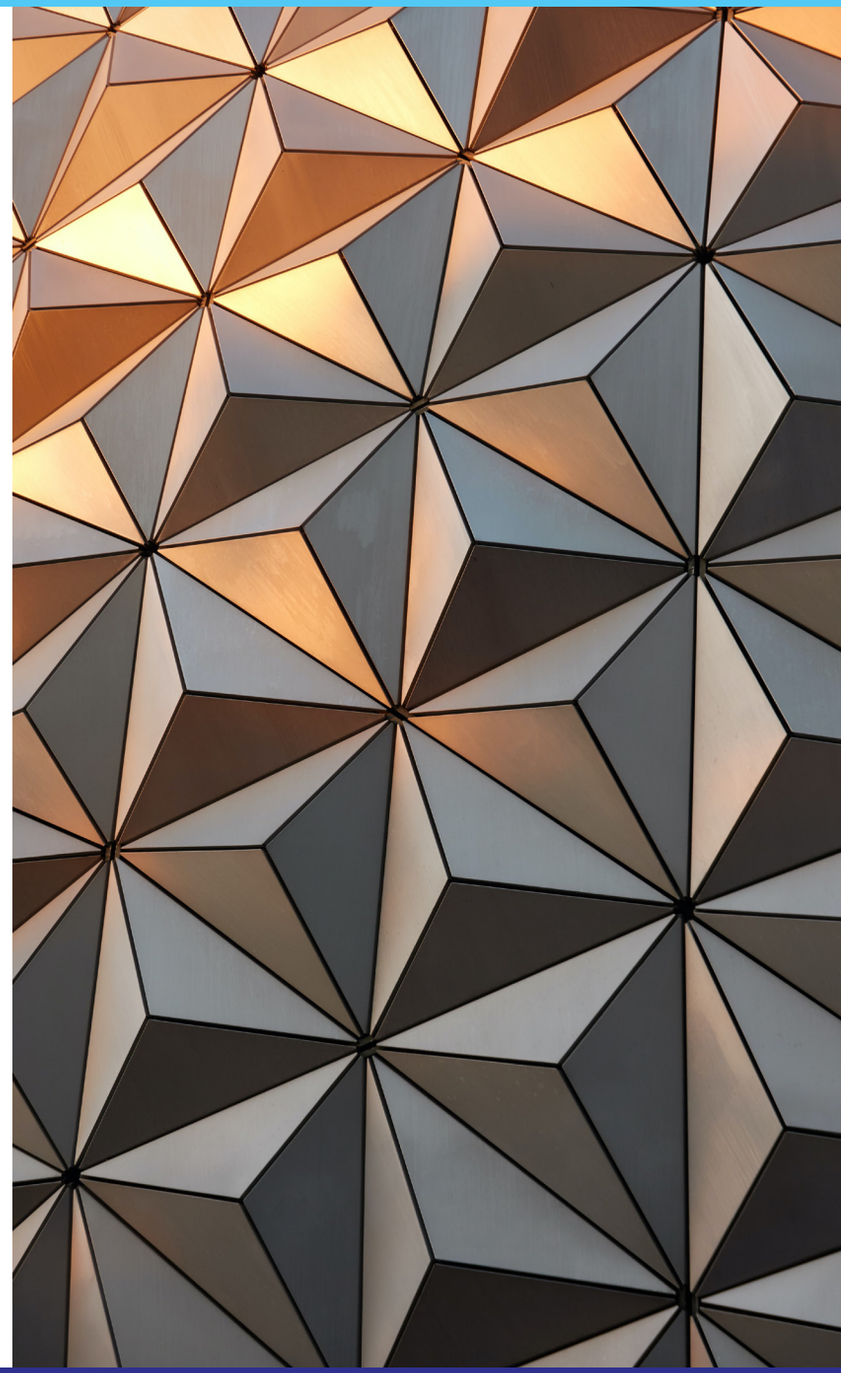

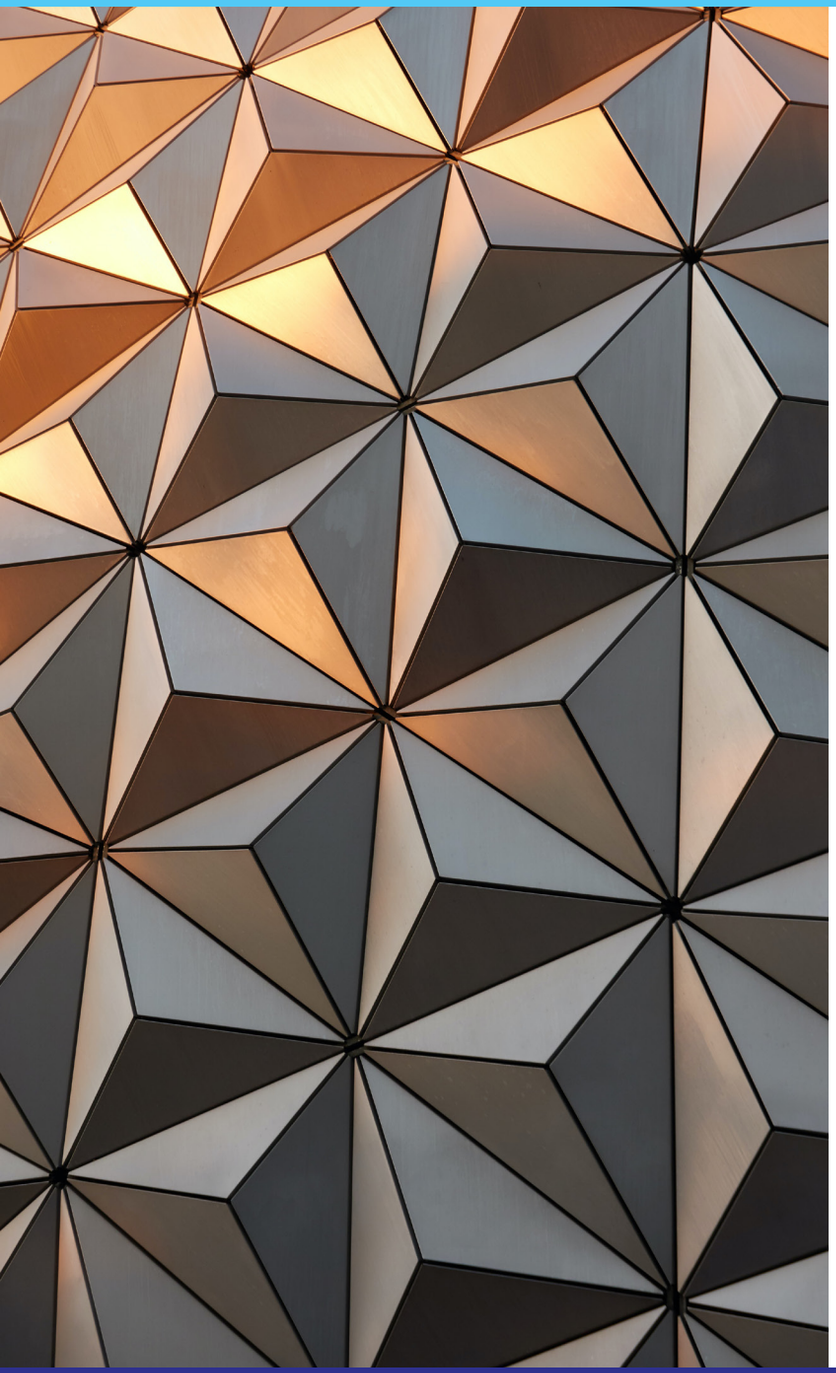

versed in both Bayesian and Frequentist methods can build a strong regulatory strategy using such an approach, but many sponsors fear the downstream effects if things were to go wrong.

Bayesian statistical approaches also enable sponsors to benefit from a number of quantitative tactics, from using predictive probabilities to estimate the likelihood of early stopping, to the most efficient uses of historical data. They can help quantify the bias of missing data, perform sensitivity analyses, and, with the help of simulations, make critical forecasts during a clinical trial. A number of these approaches can be implemented regardless of whether the clinical trial is designed using Bayesian or Frequentist statistics.

The use of Bayesian statistics should not be confused with the use of Bayesian decision-rules for clinical development strategy. In the case of clinical development strategy, Bayesian approaches offer key insights. Essentially, Bayesian decision-rules allow sponsors to take clinical data and forecast expected utilities of different strategic options, to create a statistical design that optimizes strategic goals. They can be used to select primary endpoints for clinical trials, or to develop Go/No-Go rules in early-phase clinical trials. They can also help sponsors to streamline numerous considerations that can arise during a clinical trial from concerns about cost, time, and sample size, to anxieties about treatment effect.

Read the original post.

Cytel