

# The Value Proposition

**Tom Parke at Tessella demonstrates the benefits that adaptive clinical trials can provide, and argues that although they may require greater initial effort, the end result is a more efficient development programme**

Pharmaceutical and biotech companies are increasingly interested in the use of ‘adaptive’ clinical trials to improve the process of drug development. By an ‘adaptive trial’, we mean a trial that is designed from the outset to review key response data as we collect it, and carry out certain pre-specified adaptations in the light of what we learn.

In large confirmatory trials, regulatory concerns will limit adaptations that can be considered to actions such as stopping the trial early, performing a sample size re-assessment or dropping a treatment arm. In early phase trials, frequent adaptation and a wider scope for adaptation are permitted, for instance: frequent optimisation of the probabilities of adaptation to the different treatment arms; more frequent interims to choose to drop or add treatment arms or to stop the trial; modification of inclusion and exclusion criteria to drop patient sub-populations from the study or restrict the study to a sub-population.

These kinds of trials are more complex, take longer to design and require a more dynamic approach to running clinical trials, but there is already good evidence that these efforts can reap great rewards. The advantages of adaptive clinical trial designs are quicker decision-making, better information at the end of the trial, greater flexibility in the trial design, and the ability to optimise the overall development programme for a compound.

## FASTER DECISION MAKING

The simplest way that adaptive trials achieve faster decisions is through the ability to stop the trial early, saving time, resources and money. In an adaptive trial, the key trial data is collected quickly, and monitored regularly and frequently. This is usually done using an automated analysis that generates a report for review by the data monitoring committee. The frequency of these reviews depends on the speed at which the data accrues, but in Phase II dose finding trials with an expected length of six to 18 months, these are typically weekly or fortnightly. The review of the

resulting report usually takes only a few minutes, until the weight of evidence supporting a decision becomes overwhelming and longer reviews are required to confirm that this is what the company wants to do. These reviews may take into account factors not included in the adaptive algorithm, such as market conditions and other response data.

The decisions that can be taken include stopping the trial for futility or efficacy and making a ‘phase transition’ within the trial. For instance, sufficient evidence of drug effect at the highest dose may allow us to open up lower doses to see if they too might be effective – transitioning from a Phase IIa ‘proof-of-concept’ to a Phase IIb ‘dose finding’ stage. Alternatively, sufficient evidence that we have located a tolerated dose that is sufficiently better than placebo may allow us to drop all but the selected dose – transitioning from a Phase IIb ‘dose finding’ stage to a Phase III ‘confirmatory’ stage. Lastly, decisions can be made that do not have an impact on the trial but affect the development programme at large – for instance, to invest in further manufacture of the compound or to start a parallel trial using the development compound in conjunction with an existing therapy or on an alternative indication.

Even when the decision to be taken is a ‘one-off’, such as dropping a treatment arm, there is huge value in monitoring frequently whether one wants to make this decision. Designs that adapt using a small number of interim looks at the data suffer from two serious problems. Firstly, at the point at which the interim is taken (after a fixed time or fixed number of subjects have been recruited) there may not be sufficient information to make the

decision. Secondly, it may take so long to collate and clean the data that by the time the interim decision can be made, it has been overtaken by events – there are cases where the trial had finished recruitment before the interim decision could be made.

Out of any drug development pipeline, the majority of compounds will fail (2003 figures give ~60 per cent failing at Phase II and ~30 per cent of those remaining failing at Phase III. Current trends in the industry suggest that today these figures are worse) (1). Adaptive trials are particularly good at identifying failing compounds. Fitting a dose response curve to the data results in some smoothing of the estimate of response at adjacent doses, and favouring randomisation to doses where some response is seen.

Figure 1 shows a simulation of an adaptive design to illustrate this. In this graph, the fine green line is the mean dose response that is being simulated, in this instance, the ‘null hypothesis’ that the drug performs no better than placebo. The black circles show the means of the simulated patient responses, the heavy green line the estimated dose response curve fitted to the data with the vertical green bars at each dose showing the credible interval of the estimate. At the bottom of the graph, the blue bars show the proportion of subjects allocated to each dose.

Notice how the fitted dose response curve avoids rising to the high means of the observed responses at doses three and five. This shows the effect of the smoothing from the dose response model that is being fitted to the data. More subjects have been allocated to these doses than the neighbouring doses where the response is weaker. If the trial continued with more

data on these doses, we would expect the observed mean response to more closely resemble the true underlying mean. This particular simulation stopped for futility after 133 subjects were recruited from a possible maximum of 250. We have seen companies estimate that they made net savings of \$3.5 million and \$1.5 million in adaptive trials due to stopping trials early.

Trials do not always fail because the compound does not work. They can fail because of problems with the protocol (including the wrong subjects), a poor formulation (compound not reaching the active site in the body), or with the endpoint (insensitivity, a huge placebo response or larger than expected variability). Only by analysing the data at an early stage during the trial can these problems be picked up soon enough for the trial to be stopped with enough time and budget left to run a trial with the problem fixed. Some people shy away from adaptive trials because they feel, at

the moment, that they are too risky, although arguably adaptive trials allow the trial risks to be better managed.

### BETTER INFORMATION

Adaptive trials deliver better information by adjusting the trial as it runs, typically in altering the proportions in which subjects are randomised to treatment arms or by dropping unpromising arms altogether. This results in a greater proportion of the subjects being treated in the treatment arms that are most likely to be carried forward to the next trial. In addition, by fitting models to the data, the conclusions drawn can be made with greater confidence and we can learn more from the data.

In order to attain an estimate of how much more useful the information from an adaptive trial is compared to a conventional one, I have run some simulations of the simple example trial devised by the PhRMA Adaptive Dose-

Ranging Studies Working Group (2). An adaptive NDLM design in a trial using 250 subjects over nine treatment arms (placebo plus eight doses of the study drug) in a set of simulated trials, on average allocates 32 subjects to the dose that was, at the conclusion of the trial, judged to be the minimum efficacious dose (MED). By design, one third of the 250 subjects are allocated to the placebo arm, so 19 per cent of the 167 subjects that have been allocated to doses of the study drug have been allocated to the MED. Whilst in a similar, but non-adaptive trial, 12.5 per cent of the patients allocated to doses of the study drug would be allocated to the MED. Thus the equivalent non-adaptive study would require 339 subjects to achieve the same final number of patients on the MED. Furthermore, because of the modelling of the dose response, the accuracy of the estimate of the response at the MED is better than in a conventional study. The average standard deviation (SD) of the estimate of the mean response from the fitted model is 0.30, whereas the average SD of the estimate of the mean response just from the data available at that dose is 0.36. This may not seem like a large difference, but to obtain the same level of accuracy from a non-adaptive design without modelling would require over 500 subjects. Thus, in this instance, a conventional trial would have had to have been twice the size to yield the same information as the adaptive trial.

Figure 1: Allocated subjects and fitted dose response for a simulated ineffective compound

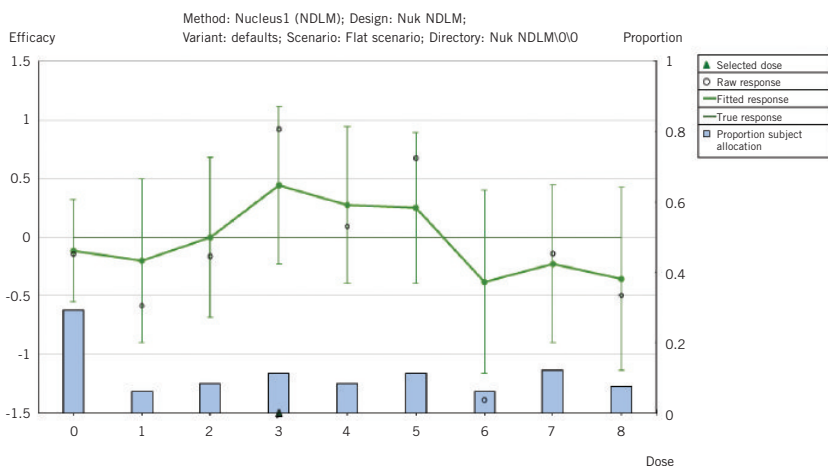


Figure 2: Allocated subjects and fitted dose response for a simulated effective compound

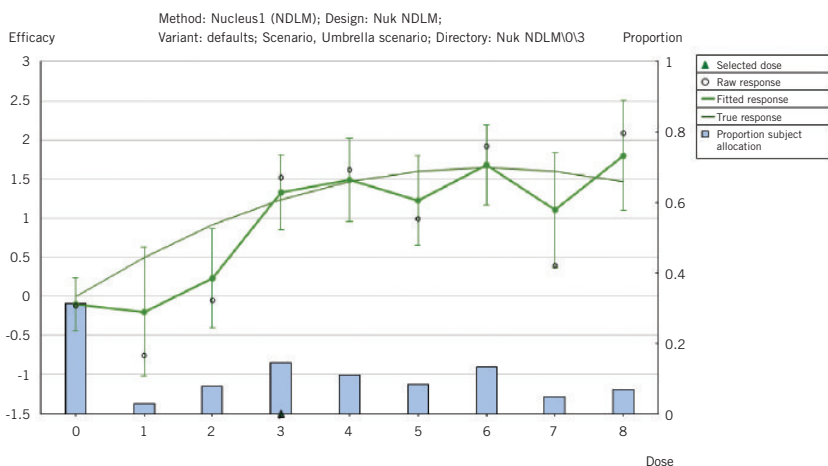


Figure 2 shows the response data at the end of one of the simulate trials. The graphical conventions are the same as for Figure 1. Note how the proportion of subjects allocated to dose 3 (the dose selected as the minimum effective dose) is relatively greater than for most of the other doses. Also, note how the credible interval bars are shorter around the estimate of the mean response for dose 3, than say dose 1 or 2. We have more confidence in the estimate of the mean response at this dose than we would have after a conventional trial because a higher proportion of subjects have been allocated to it by the adaptation, and because of the curve fitting (note how similar the response is at the adjacent dose).

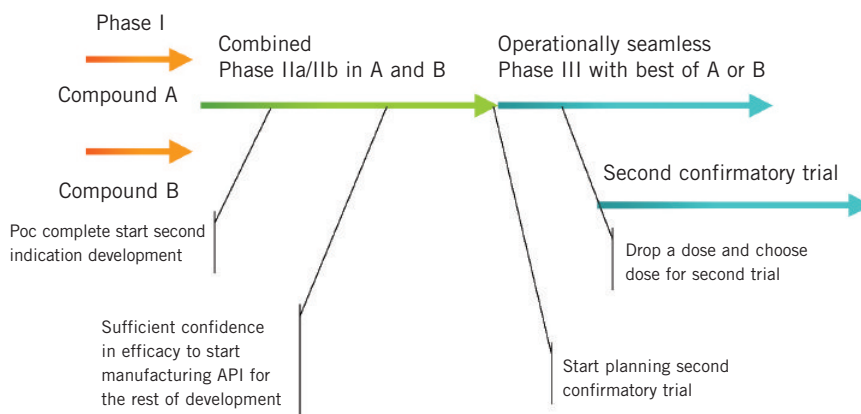
### GREATER FLEXIBILITY

Adaptive trials allow greater flexibility in the specification of the goal of the trial,

and more opportunities for implementing decisions as the trial runs. This not only allows different conventional trials to be combined in a single adaptive trial as discussed above, but new types of trial can be considered. For instance, often a pharmaceutical company will have both a 'lead compound' and a 'backup compound' as candidates to treat a particular disease. It may be worth considering testing both compounds in a Phase II adaptive trial that adapts the allocation to favour not only the most promising dose, but the most promising compound. This saves considerable costs compared to running two separate trials and saves time by avoiding testing the lead compound, having it fail and then switching development to the backup compound.

A very different approach is necessary for compounds in therapeutic areas where the same compound could treat a number of different diseases – for instance different types of cancer tumour. Here there is value to be gained from running parallel studies in all the different diseases and modelling the drug's likely effectiveness across them all using a hierarchical model. This sort of model would take into account whether

Figure 3: Example Adaptive Development Programme



the responses are consistent across the different diseases. If there is consistency, fewer subjects will be required overall to reach the same level of confidence in the outcome, compared to analysing the results separately.

Another area of increasing interest is the identification of significant sub-populations of patients, identified by a possibly significant genetic marker, for instance. In the sub-population, the drug may be particularly effective or ineffective, or particularly tolerable or toxic. Identification of these significant sub-populations in some areas is becoming increasingly important for the successful development of drugs. A particularly ambitious and innovative design along these lines is the proposed 'I-SPY 2' trial in neoadjuvant breast cancer (3). The value of such opportunities is hard to put a figure on. They may range from a 20 per cent saving in direct trial costs, to making the identification of successful drug possible.

#### MORE EFFICIENT DRUG DEVELOPMENT

Adapting early phase trials prompts us to do what was almost unthinkable a decade ago – to look early and often at the data being collected during a trial. The data can be kept from the team running the trial, the investigators and subjects, but used to make development decisions earlier – before the trial is complete, for instance:

- Are the results good enough to justify manufacturing more of the compound in order to run further trials shortly after this one is complete?
- Do the results warrant starting to plan the next trial before this one is complete?

- Do we have sufficient safety data to start a second trial with this compound in a second indication or different population?
- Is the compound failing and should we be advancing the backup compound?
- Does the biomarker data justify the use of the biomarker on the next trial?

All these decisions can be made earlier, or with better information, because we are not waiting until the end of the trial to look at the data. Indeed, even without adaptation, continuous monitoring and modelling of the endpoint data can be hugely beneficial. Once we realise this, these additional or supplementary questions to aid the rest of the development programme can be considered as targets for the trial design and real-time trial analysis.

Adaptive clinical trials are not just better ways of running some types of clinical trial; they offer the prospect of smarter, more intelligent, more efficient development programmes.

#### References

1. DiMasi JA, Hansen RW and Grabowski HG, The price of innovation: new estimates of drug development costs, *Journal of Health Economics* 22: pp151-185, 2003
2. <http://biopharmnet.editme.com/PharmaAdrsHome>
3. <http://gforge.nci.nih.gov/docman/view.php/559/15128/1%20SPY%201%20and%202%20overview%20v2-%20MBuxton,%20SMadhavan,%20LEsserman.ppt>

#### About the author



Tom Parke joined Tessella in 1996, and has worked on supporting adaptive clinical trials since 1998. He managed the development of the system that allowed Pfizer to run their ground-breaking ASTIN trial, and has worked on numerous further dose response adaptive trials. He has overseen the development of algorithms, simulation tools and systems to support adaptive trials, as well as integrating adaptive algorithms with existing central randomisation and electronic data capture systems. Whilst currently working on his twelfth dose-response adaptive trial, Tom is consulting with a number of international pharmaceutical companies, helping them to define the software systems they require to run adaptive clinical trials.  
Email: [tom.parke@tessella.com](mailto:tom.parke@tessella.com)