STATISTICS IN MEDICINE
Statist. Med. **18**, 3425–3434 (1999)

ON BEING THE STATISTICIAN ON A DATA AND SAFETY MONITORING BOARD†

JOHN WHITEHEAD*

Medical and Pharmaceutical Statistics Research Unit, The University of Reading, P.O. Box 240, Earley Gate, Reading RG6 6FN, U.K.

SUMMARY

Data and Safety Monitoring Boards (DSMBs) are commonly appointed to monitor emerging data from major clinical trials. This paper describes their organization and remit, and their relationship with other trial committees and structures. The operation of formal stopping rules for safety and efficacy by a DSMB is discussed. The duties of a DSMB, from pre-trial planning through to stopping a study are described in detail, with emphasis on the reporting of information to the DSMB and the reporting of conclusions by the DSMB. The issue of blindness is given prominence and the role of the statistician on the DSMB is explored in detail.

1. INTRODUCTION

Data and Safety Monitoring Boards (DSMBs) have become a familiar feature of major long term clinical trials in life-threatening diseases. They are regularly used in trials sponsored by the pharmaceutical industry and are becoming a common part of the procedures of public sector trials run by bodies such as the MRC in the U.K. and the NIH in the U.S.A. The names of these committees vary; most make some reference to 'data', 'safety' or 'monitoring', but less explicit titles such as 'Independent Review Panel' are sometimes adopted.

The activities of DSMBs were described by Armitage,[1] and a whole issue of *Statistics in Medicine* was devoted to them in 1993.[2] An informative summary of their role has been given by DeMets.[3] The constitution of a DSMB was discussed by Armstrong and Furberg.[4] Regulatory and public sector guidelines[5,6] also describe such boards. This paper presents a personal view, gained from service on and reporting to more than 20 DSMBs in a variety of therapeutic areas. The trials concerned have involved both pharmaceutical and public sector sponsors, have ranged from small phase II studies to large definitive studies with long-term follow-up, and have varied considerably in both design and administration. Owing to my own research interests, an unrepresentatively large proportion have involved some type of formal sequential procedure. My own role has been exclusively statistical.

* Correspondence to: John Whitehead, MPS Research Unit, The University of Reading, PO Box 240, Earley Gate, Reading RG6 6FN, U.K. E-mail: J.R.Whitehead@reading.ac.uk.

† Presented at the International Society for Clinical Biostatistics, Nineteenth International Meeting, Dundee, Scotland, August 1998

CCC 0277–6715/99/243425–10$17.50

No single model can cover all DSMBs. In a major trial in a life-threatening disease, recruiting thousands of patients and treating them over several years, a large DSMB with a complex structure of administration and reporting might be appropriate. In a shorter-term phase II study with fewer than 100 patients, such elaborate procedures would be out of proportion, and a simplified approach would be more reasonable.

The structure and remit of DSMBs will be described in the following section. Section 3 concerns the use of formal stopping rules based on safety considerations, and the role which a DSMB can play in operating an efficacy stopping rule. In Section 4 the operation of a DSMB is presented in detail, from pre-trial planning through to procedures for stopping the study when the need arises. The form of reports to the DSMB and the question of blinding these are considered in detail. Section 5 draws conclusions, particularly for statisticians who serve or might be called upon to serve on a DSMB.

Throughout the paper, and especially in Section 4, personal views are expressed. Implicit in the advice offered is the phrase, 'in the author's opinion'. It is also inevitable, given the wide range of trials which employ a DSMB, that there will be exceptions to almost every rule.

2. DATA AND SAFETY MONITORING BOARDS

As a clinical trial proceeds, data concerning the relative merits and relative dangers of the treatments under study accumulate. In pharmaceutical trials, the company's pharmacovigilance department will receive reports of deaths and serious adverse events, and usually these will also be seen by regulatory bodies. However, it is normal for these reports not to identify the treatment received. The mechanism allows a reaction to an excessive overall rate of serious adverses events or mortality, but not to an excess of events on the experimental treatment relative to the control. It is generally regarded as essential that comparative trial data be kept from the investigators and the sponsor until the study has been completed,[3,5–7] so as to avoid biases in its conduct. (However, for a contrary view, see Meinert.[8]) It is to allow monitoring of comparative trial data during the study, without compromising the investigators or the sponsor, that most DSMBs are set up. For the purpose of this paper, a DSMB will be defined as those people who are privy to comparative trial data during the course of its conduct, and who can make recommendations to modify or stop the trial as a result.

Some trials, such as those comparing surgery with best supportive care, cannot be blinded. The need for a DSMB in such cases is less pronounced, as their function could be subsumed by the Steering Committee. However, it may be preferred that formal tabulations by treatment do not receive wide circulation, and a DSMB might thus be used to prevent this. Whilst not exactly blind, other parties to the trial will only have a restricted view of the comparative picture.

A DSMB will comprise clinical experts in the condition being treated, and a statistician whose role is to put the emerging results into the context of what might happen by chance alone. In some cases, specific side-effects are anticipated and clinicians specializing in these conditions are also included. Some DSMBs have a lay person as a member – perhaps a patient representative or an ethicist. Such DSMB members help to ensure that 'common sense' considerations are not overlooked amidst the clinical and statistical detail under debate. The DSMB should have a chairperson, and a secretary to take minutes and deal with correspondence. The latter could also be a DSMB member, or a professional secretary, depending on the scale of the task. The minimum size is two clinicians and one statistician; clinicians are rightly uncomfortable about serving in isolation. In larger trials, DSMBs may have as many as ten members. Members of the

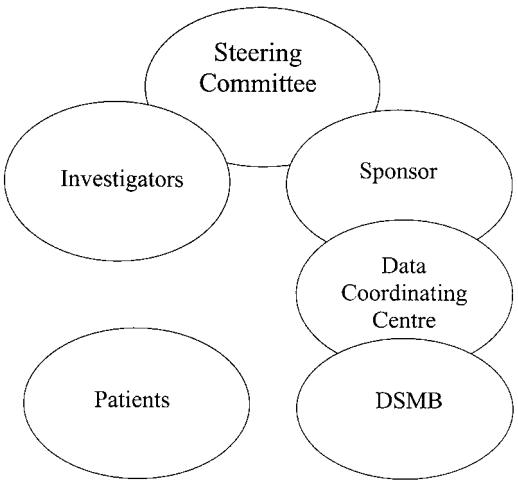

Figure 1. Interrelationships between the parties in a large clinical trial

DSMB should have no financial stake in the success or failure of the trial; in particular they should not deal in the shares of the sponsor during the conduct of the trial.

The remit of the DSMB is quite simply to protect the ethical and safety interests of the patients. In particular they are to do this using their privileged access to comparative trial data.

Before proceeding to discuss the role of the DSMB further, it is helpful to clarify the meaning of blindness. A person will be said to be *totally blind* if they do not know what treatment group any patient is in, or whether any two patients are in the same treatment group. Such a person has no access to any listings or tabulations broken down by treatment group, even with these coded as A and B (say) rather than being identified by name. Following terminology used in reference 9, a person will be said to be *subgroup unblind* if they know whether each patient is on A or B, but they do not know what A or B are. This term will be used to indicate anyone who sees coded listings or tabulations by treatment group. A person will be said to be *totally unblind* if they know explicitly what treatment group every patient is in, or if they see listings or tabulations labelled by treatment name.

The structure of a large clinical trial is illustrated in Figure 1. Each oval represents a party to the trial, and direct overlaps indicate permissible common membership. Every trial (large or small) should have a clearly identified Steering Committee, comprising some of the investigators, representatives of the sponsor and independent experts. It is the responsibility of the Steering Committee to prepare the protocol, to ensure that the trial follows it, and to institute protocol modifications when these become necessary. The trial sponsor may be a pharmaceutical company, a public sector body such as the MRC or the NIH, or a medical charity. No investigators and no representative of the sponsor should be a member of the DSMB.

In between the sponsor and the DSMB lies the Data Co-ordinating Centre which will receive data from the study centres, enter them into a computer database, and prepare reports for the DSMB. Often these roles are separated. Data entry may be performed by employees of the sponsor or by an independent contract research organization or similar body. All data except the

treatment allocation can be handled in this way. Reports on comparative data for the DSMB will be prepared by one or more subgroup-unblind statisticians, from the data file, plus a file identifying the treatment associated with each patient by codes such as A and B.

The comparative reports may be prepared by the statistician on the DSMB. Time has to be allowed for this process; it is unrealistic to expect it to happen at the DSMB meeting itself. Usually, the DSMB statistician will require assistance, preferably from just one member of his or her own staff who is also independent from the investigators, Steering Committee and sponsor, and who will become subgroup-unblind. Alternatively, the reports may be prepared by a subgroup-unblind statistician who is an employee of the independent contract research organization, working in collaboration with the DSMB statistician. A third option is for reports to be prepared by a specific subgroup-unblind statistician who is a member of the sponsor's staff, charged to keep the information seen confidential from his or her employer and colleagues. This arrangement can work satisfactorily provided that all involved act responsibly. It is understandable in the context of small-scale trials, but is probably best avoided in high-profile pivotal studies. The 'Rolls-Royce' solution is for one independent subgroup-unblind statistician or small statistical group to receive totally blinded data from the Data Co-ordinating Centre and to report to a DSMB which has its own statistician, although such an elaborate procedure may be too expensive for all but the most major trials.

3. FORMAL STOPPING RULES

In order to help fulfil its remit in protecting the safety interests of patients in the trial, a DSMB might use a formal stopping rule for safety. As a separate issue, the DSMB might be requested to help in the operation of a formal stopping rule for efficacy.

A formal safety rule is helpful in trials in which one form of event is of greatest concern. This may be a specific side-effect, such as liver damage, or it may be early mortality when in most respects this summarizes other safety concerns. The construction of safety rules and their operation in some recently completed trials is the subject of a separate paper. Usually a binary response is defined, such as death within two weeks, and a statistic contrasting the event rates on experimental and control is calculated at each interim safety analysis. The formal rule will indicate stopping if the event rate is sufficiently greater on experimental than on control. The rule is chosen to have acceptable stopping probabilities when there is in fact no true safety difference, and when a safety problem of a given magnitude exists.

Formal safety rules have several advantages. They ensure that stopping scenarios are discussed thoughtfully by the DSMB while they are being devised. They operate to specified error probabilities. They provide a consistent criterion of judgement each time the DSMB meets. However, they should be used in addition to the data presentations normally considered by a DSMB. It may be that, although the form of event chosen to reflect safety concerns is equally common in the two treatment groups, these presentations identify other unanticipated dangers of the treatment. Conversely, an excess of the event used in formal monitoring might be counterbalanced by an equally dramatic beneficial effect leading to a desire and an ethical basis for continuing the study. Consequently, formal safety rules should be overridden, in either direction, if clinical judgement deems it necessary.

The response used in operating a safety rule should normally be one that is available early during a patient's course of treatment, so that a quick reaction to imbalance is possible. It may be correlated with the response used to assess efficacy, but will not usually be identical to it. The

operation of a safety stopping rule will have an effect on a final frequentist efficacy analysis of the trial. Usually this is negligible in magnitude, and the conventional efficacy analysis will be made conservative, that is it will understate significance by reporting too large a p-value. The conservatism arises because treatments can only be rejected by the rule, not recommended, so the probability of claiming significantly beneficial efficacy is (slightly) reduced under all hypotheses including the null. The power is very slightly reduced too. In equivalence trials, the probability of claiming equivalence is also reduced, as no claim will be made after stopping for a safety concern. Corrections for the safety stopping rule when performing the efficacy analysis are possible using bivariate sequential methods;[10–12] they will lower the p-value but are unlikely to have an appreciable effect.

Stopping rules for efficacy are used to allow a conclusion to be drawn concerning the main trial question, as soon as sufficient data are available. Extensive literature on various forms of sequential design for this purpose is available.[9,13–15] Whereas a safety stopping rule is of prime concern to the DSMB, and will often be devised by them, an efficacy rule is more likely to be proposed by the Streeing Committee or the sponsor. Stopping for a positive efficacy difference may be to the advantage of both future patients and the sponsor. Stopping early in the case of lack of effect or inferiority of the experimental can save resources and allow patients to be entered into more promising studies. There are both ethical and economic aspects which extend beyond the patients recruited into the trial. Although scientific interests may be served by additional data collected in a trial already clearly indicating a lack of treatment effect, consideration has to be given to whether the sponsor would wish to continue paying for the study and whether patients would continue to consent to be randomized, were this lack of effect generally known.

In operating a sequential efficacy design, the DSMB acts as a buffer between the narrow computer-generated finding and the irrevocable decision to inform the Steering Committee that a boundary has been crossed. The DSMB can override the formal rule. The rule should be discussed and accepted by the DSMB before the trial begins, and so it is likely only to be overridden if the assumptions underlying its creation are suspected to be false. These assumptions include homogeneity of treatment difference over subgroups, validity of models such as proportional hazards and an equal or favourable balance of side-effects. When the validity of such assumptions are in doubt, it might be prudent to collect more data to ensure a robust analysis of sufficient power. Sometimes all assumptions may appear valid, but what appeared to be a sensible design before the trial began might seem to be less sensible when the stopping criterion is reached. Current judgement must then override the design, but the initial misjudgement must be acknowledged and appropriate lessons learnt for future trials. Because of the effects that departure from a predetermined sequential design can have on an eventual frequentist analysis, DSMBs should accept them as firmer guidelines than safety stopping rules, being aware of the negative effect that casually ignoring the design would have on the trial's credibility, and recording carefully any reasons for overriding the efficacy stopping rule.

4. OPERATION OF A DSMB

4.1. Meetings and reports

A DSMB should be appointed before the trial begins. This allows them to meet, in person, with representatives of the Steering Committee and the sponsor to learn about the trial and to review the protocol while there remains a chance to make changes. In reviewing the protocol, the DSMB

Table I. Statement of the terms of reference for a typical Data and Safety Monitoring Board

Remit:

To protect the ethical and safety interests of the patients

Duties:

1. To receive periodic reports of data on adverse events, laboratory findings and efficacy measures, and also individual adverse event forms.
2. To consider whether such data are consistent with the ethical and safe continuation of the study, paying attention to any predefined guidelines or stopping rules.
3. To consider whether the data received are sufficient and are reported quickly enough for fulfilment of their obligations.
4. To document their major deliberations, and in particular any reasons for departing from guidelines. These minutes may be requested by the sponsor or a regulatory body at the end of the trial.
5. To issue periodic reports giving one of the recommendations below:
   (a) The study should continue without modification.
   (b) The study should continue with the following modifications ....
   (c) The study should be stopped due to the following safety concerns ....
   Patients currently under treatment should ....
   (d) The study should be stopped, as a result of the formal stopping rule for efficacy.
6. To keep all information received confidential, and especially to keep comparative information from the investigators, the Steering Committee and the sponsor.
7. The DSMB can consider and comment on matters such as recruitment rate and overall event rate. However, the Steering Committee is also free, and perhaps more appropriate, to consider these issues.

needs to be comfortable about the purpose, methodology and ethical aspects of the trial, but should accept that responsibility for detail rests with the Steering Committee. The DSMB should not try or be asked to play both roles. The operation of any efficacy stopping rule should be considered, and the DSMB needs to decide whether it should adopt or devise a formal safety stopping rule. The implications of any stopping rules need to be thought out carefully in advance, and the circumstances in which they would call for termination or continuation confirmed as being intuitive and ethical. Formal terms of reference, such as those listed in Table I, should be prepared.

During the trial, the DSMB will receive various reports. Sometimes it is appropriate for them to be sent individual serious adverse event forms as the events occur. It can be arranged for these to be routed through the subgroup-unblind statistician, who can attach treatment labels so that emerging imbalances can easily be seen. At regular intervals, full reports will be sent to the DSMB, comprising adverse events and laboratory data by treatment, demographic and baseline data to put the former into perspective, and presentation of any formal safety or efficacy interim analyses. After scrutiny of each full report, the DSMB will discuss their conclusion and issue a brief report to the sponsor stating whether the study should continue or stop. In the case of continuation, protocol amendments might be suggested; in the case of stopping, reasons might be given. The discussions of the DSMB might be face-to-face or by teleconference. It is advisable for a face-to-face meeting to occur at least once every six months to ensure that each DSMB member has a full opportunity to express his or her views. Intermediate meetings (if needed) might take place by telephone.

Representatives of the sponsor or the Steering Committee might attend an open part of these meetings in order to present information on the state of the trial, and to receive the report of the DSMB. Discussion of the safety report and of the decision to recommend continuation or stopping will, however, take place in a closed session of the DSMB alone. Alternative arrangements include having the sponsor's representative available to answer questions by telephone. Minutes of each meeting need to be kept, to be made available to the sponsor, and if requested to regulatory bodies, after the trial has been completed.

4.2. The form of reports submitted to the DSMB

Reports sent to the DSMB are comprehensive and lengthy. As they contain important information which has to be assimilated quickly during a short period between being received and the DSMB meeting, careful planning is needed to ensure that they are as clear and concise as possible. This planning is the responsibility of the Data Coordinating Centre, with input from the DSMB, and needs to take place well in advance of the production of the first report. The DSMB's role will mainly consist of stating what they want to see; it is the Data Coordinating Centre which will determine how it is presented. Of course, revisions to the content and format of the reports will be made as the study progresses.

Detailed guidance concerning the construction of reports for a DSMB lies outside the scope of this paper. The considerations of layout and presentation are those which underlie all statistical reports, but those prepared for DSMBs often appear to be little more than a 'dump' of the data available. Attention should be given to the ordering of tables, to placing all information relevant to a single issue in close proximity and to avoidance of duplication of information. Percentages should be used with care, and chosen to make relevant comparisons easier. Graphical presentations should be used to accompany key tables, but should not be used so frequently as to lessen their impact. It is best to consider what is wanted first, and then to find a way of forcing the computing software to produce it, rather than being dictated to by the conventions of statistical packages. It is never easy to judge clarity by table shells. The statistician responsible for the report needs to study pre-release runs with real data and to go through the exercise of making sense of it, before running it for submission to the DSMB.

In addition to receiving reports on the trial being monitored, the DSMB may also receive reports on other related trials. Before the trial starts, the DSMB should be briefed by the sponsor on the outcome of completed trials of the same or similar interventions. If related trials are in progress simultaneously, then news of those is also relevant. Various arrangements are possible. A sponsor might choose to ask a single DSMB to oversee all of their trials of a particular intervention. Alternatively, separate DSMBs might be set up, but they might exchange the minutes of their meetings, thereby advising one another of any concerns. Usually the result will be mutual reassurance. It is advisable for each trial to be run to its intended completion (whether governed by a fixed sample size or a sequential stopping rule) if positive efficacy results are being found. In this way, each trial fulfils its own power specification and can serve as independent corroboration of the other(s). When similar safety concerns are found in more than one study, then less formal combination of results can be appropriate; it may be sufficient to seek evidence of safety problems from pooled or stratified data.

In the Bayesian approach to sequential clinical trials,[14] incorporation of data from an external trial can be achieved in a formal way. This is perhaps more appropriate when judging safety

concerns than for combining efficacy results. Sequential designs allowing repeated meta-analyses can also be constructed.[16]

4.3. Blindness

The reports to the DSMB should be totally unblind, identifying treatments by name. This is because, to fulfil their role of protecting the safety interests of patients, the DSMB should have access to all of the facts. In practice, data are often presented with treatments coded as A and B. While clinicians often seem to favour this approach, statisticians appear to be less enthusiastic.[17] Even when coded presentation is used, it is an overstatement to refer to the procedure as 'a blinded review'. First, the treatment identities are often revealed through the pattern of side-effects. Once this happens, no purpose is served in maintaining the pretence of blindness; indeed it can be dangerous because there is a small chance that the treatment identities have been wrongly guessed and yet these false identities are influencing attitudes towards trial continuation. Elaborate methods aimed at preventing unblinding are sometimes devised, in which some listings are coded A and B while others receive the designation C and D. Often features such as treatment group totals allow this form of blindness to be broken, but as its intention is to deny the DSMB a complete picture of the interim data it should not even be attempted. Second, it is usual for the DSMB chair to have a sealed envelope containing the identities of treatments A and B. The envelope can be opened whenever the DSMB see fit, and is likely to be opened prior to stopping the study (unless the stopping is for lack of effect). This feature alone should prevent the use of the term 'blinded review' in any description of the trial procedure.

4.4. Unexpected imbalances

Every so often, interim data reveal imbalances in outcome which, if they truly reflected reality, would be of serious concern. The DSMB statistician is then asked questions such as 'What is the probability of the excess on the experimental treatment of gastrointestinal disorders amongst the elderly occurring by chance?'. Of course there is no simple answer to this question which adequately allows for the number of comparisons and subgroups under review and the number of times that the data are examined.

Although there are no rigorous answers to questions such as the above, it is the role of the DSMB statistician to provide some form of response. The realm of the unexpected question can be reduced at the planning stage by identifying a primary safety endpoint and devising a formal stopping rule, as described in Section 3. Precise answers about that particular endpoint are then available. Nominal p-values can at least provide reassurance when they are large enough ($>0{\cdot}05$!); when they are smaller, subjective mental adjustment for multiple and repetitive analyses can help to form a qualitative judgement rather than provide a numerical answer. Sometimes a formal rule can be created mid-trial to help in monitoring a concern which has arisen; once more this is a qualitative rather than a precise procedure in such data-motivated circumstances. Bayesian statisticians will wish to compute posterior probabilities, but they too should avoid interpreting numerical values literally, as priors for each conceivable situation will not have been elicited in advance.

4.5. The form of reports issued by the DSMB

After each examination of the data, the DSMB will issue a short formal report to the Steering Committee. The central part of this will be a form of words similar to one of the following:

(a) The study should continue without modification.
(b) The study should continue with the following modifications ....
(c) The study should be stopped due to the following safety concerns ....
Patients currently under treatment should ....
(d) The study should be stopped as result of the formal stopping rule for efficacy.

Option (a) is the most common. The modifications in option (b) might relate to eligibility criteria, or to the collection of extra safety data. There might be modifications to encourage a closer adherence to the existing protocol. When the recommendation is to stop due to safety concerns, the DSMB will need to identify precisely what these are, summarize the evidence on which they are based, and provide recommendations about how patients currently under treatment and those who have completed treatment should be dealt with. It would be unusual for a Steering Committee not to comply with a recommendation to stop for safety, but they should be acquainted with the evidence before issuing the termination instructions.

If the efficacy stopping criterion is met, and confirmed by the DSMB, then the Steering Committee should receive the message (d). If they approve stopping, then patients currently under treatment will then have their course completed or interrupted as appropriate, and all remaining data will be collected and computerized. The database will be checked and locked. The final interim will be repeated. If stopping is confirmed, then and only then will the circumstances of the stopping be revealed. The conclusion might be that the active treatment is better than control. Alternatively, depending on design, it might be that there is no effect, or that there is a negative effect (but not to the extent that warrants stopping due to safety as in option (c)). There are advantages for the closing stages of the trial to be conducted without knowledge of what the outcome has been, and so it might be appropriate for the Steering Committee to receive *only* message (d) from the DSMB. In other situations the need to know which conclusion has been reached, perhaps to allow modification of the treatment of current patients, might outweigh the advantages of continued secrecy. The timing of breaking the blind should be determined in advance.

It may be that completion of data collection reveals that the stopping criteria had not been reached. Sometimes, re-opening the trial may be an option. More usually it will be too late. However, when trial is stopped early, the p-value will be substantially less than the preset value of α. Borderline significance will occur only if the maximum sample size is actually reached (provided that the approach of Fairbanks and Madsen,[18] adopted in reference 9, is used). The form of analysis suggested by Whitehead[19] to deal with 'underrunning', that is to analyse a trial abandoned without reaching a boundary, will lead to a less significant (larger) p-value, but one which is quite likely to remain less than α.

Other aspects of the DSMB's report to the Steering Committee may include comments on the recruitment rate, and on the form of the reports that they receive. The delay between events occurring and the DSMB receiving notification should also be considered and commented upon. A target duration between the earliest event not included in a report and the time at which the DSMB receives it can be agreed on in advance, and checked through reference to earlier reports.

5. CONCLUSIONS

As more and more clinical trials are conducted under the scrutiny of a DSMB, it will become more common for statisticians to be involved in this work. It is an important and responsible task, and one which should be approached positively and seriously.

The major concerns that the statistician should have are the following. The pre-trial planning should be done carefully and will require the time of all DSMB members. However, it is the statistician who can perhaps best foresee the dangers of reaching conclusions in an unstructured way and who has the most to gain by paying attention to design. The statistician will also have the greatest influence on the format of the reports received by the DSMB. Thought and effort at the time when these are being devised will bring benefits of clarity and efficiency at the interim reviews. The statistician must be prepared to explain the methods of analysis, and also simple concepts such as the dangers of repetition and multiplicity of data analyses.

Usually, all of the careful efforts of the DSMB will lead to a routine conclusion that the trial can continue unchanged. This is the most desirable outcome, and Board members should avoid any temptation to create problems and issues which are unimportant for the sake of being seen to be doing something.

Notwithstanding the last remark, the task of being a statistician on a DSMB is an exciting and challenging one. It is a role that more medical statisticians will be called upon to fulfil in the future and for which they need to be prepared.

REFERENCES

1. Armitage, P. 'Interim analysis in clinical trials', *Statistics in Medicine*, **10**, 925–937 (1991).
2. Ellenberg, S., Geller, N., Simon, R. and Yusuf, S. (eds). 'Practical issues in data monitoring of clinical trials', *Statistics in Medicine*, **12**, 414–616 (1993).
3. DeMets, D. L. 'Data and safety monitoring boards', in Armitage, P. and Colton, T. (eds), *The Encyclopedia of Biostatistics*, Wiley, Chichester, 1998.
4. Armstrong, P. W. and Furberg, C. D. 'Clinical trial data and safety monitoring boards: The search for a constitution', *Circulation*, **91**, 901–904 (1995).
5. International Conference on Harmonisation. 'Statistical principles for clinical trials (E9)', London, 1998.
6. Medical Research Council. *Guidelines for Good Clinical Practice*, Medical Research Council, London, 1998.
7. Facey, K. M. and Lewis, J. A. 'The management of interim analyses in drug development', *Statistics in Medicine*, **17**, 1801–1809 (1998).
8. Meinert, C. L. 'Clinical trials and treatment effects monitoring' (with discussion), *Controlled Clinical Trials*, **19**, 515–543 (1998).
9. Whitehead, J. *The Design and Analysis of Sequential Clinical Trials*, revised 2nd edn, Wiley, Chichester, 1997.
10. Jennison, C. and Turnbull, B. W. 'Group sequential tests for bivariate response: interim analysis of clinical trials with both efficacy and safety endpoints', *Biometrics*, **49**, 741–752 (1993).
11. Cook, R. J. and Farewell, V. T. 'Guidelines for monitoring efficacy and toxicity responses in clinical trials', *Biometrics*, **50**, 1146–1152 (1994).
12. Todd, S. 'Sequential designs for monitoring two endpoints in a clinical trial', *Drug Information Journal*, **33**, 417–426 (1999).
13. Jennison, C. and Turnbull, B. W. 'Statistical approaches to interim monitoring of medical trials: A review and commentary', *Statistical Science*, **5**, 229–317 (1990).
14. Spiegelhalter, D. J., Freedman, L. S. and Parmar, M. K. B. 'Bayesian approaches to randomised trials' *Journal of the Royal Statistical Society Series A*, **157**, 357–416 (1994).
15. Scharfstein, D. O., Tsiatis, A. A. and Robins, J. M. 'Semiparametric efficiency and its implication on the design and analysis of group-sequential studies', *Journal of American Statistical Association*, **92**, 1342–1350 (1997).
16. Whitehead, A. 'A prospectively planned cumulative meta-analysis applied to a series of concurrent clinical trials', *Statistics in Medicine*, **16**, 2901–2913 (1997).
17. Meinert, C. L. 'Masked monitoring in clinical trials—blind stupidity?', *New England Journal of Medicine*, **338**, 1381–1382 (1998).
18. Fairbanks, K. and Madsen, R. 'P values for tests using a repeated significance test design', *Biometrika*, **69**, 69–74 (1982).
19. Whitehead, J. 'Overrunning and underrunning in sequential clinical trials', *Controlled Clinical Trials*, **13**, 106–121 (1992).

