

NCBI Bookshelf. A service of the National Library of Medicine, National Institutes of Health.

Smith PG, Morrow RH, Ross DA, editors. *Field Trials of Health Interventions: A Toolbox*. 3rd edition. Oxford (UK): OUP Oxford; 2015 Jun 1.

Chapter 11 Randomization, blinding, and coding

1. Introduction to randomization, blinding, and coding

As discussed in Chapter 4, the random allocation of participants in a trial to the different interventions being compared is of fundamental importance in the design of investigations that are conducted to produce the highest-quality evidence of any differences in the effects of the interventions. Only if the units to which the interventions are applied (for example, individuals, households, or communities) are randomized between the interventions under study and the study is of a sufficient size is it possible to be confident that differences in the outcome measures of the trial among those in the different intervention groups are due to the effects of the interventions, rather than to underlying differences between the groups. Randomization should ensure that any potential confounding factors, whether known or unknown, are similarly distributed in each of the intervention groups and therefore cannot bias the comparisons of outcome measures between the groups.

Randomization, if done properly, eliminates the possibility of subjective influence in the assignment of individuals to the different intervention groups. Sometimes ‘pseudo-randomization’ methods are employed in trials for reasons of convenience such as alternate assignment of the different interventions to successive trial entrants or allocation based upon the date of birth or date of entry (with, say, one intervention being assigned to those reporting on even dates and another to those reporting on odd dates). However, proper randomization is superior to any systematic method of allocation, and these other methods should be avoided, unless there are very compelling reasons for using them. With systematic allocation, it is possible for the investigator, and sometimes the participant, to know in advance the group to which a participant will be allocated, and this may introduce conscious or unconscious bias into the allocation procedure. For example, such knowledge may affect the investigator’s judgement as to whether or not an individual is eligible for entry into a particular trial. For this reason, it is essential that the randomization is done (or the randomization allocation is revealed to the investigator) only *after* it has been ascertained both that an individual is eligible for entry into a trial and also that he or she is prepared to participate in the trial, no matter which intervention is assigned.

As (Schulz 1995) pointed out, the success of randomization depends on two interrelated processes. The first entails generating a sequence by which the participants in a trial are allocated between intervention groups. To ensure unpredictability of that allocation sequence, it should be generated by a random process. The second process *allocation concealment* shields those involved in a trial from knowing upcoming assignments in advance, so that investigators cannot change who gets the next assignment, potentially making the comparison groups less equivalent and thus biasing the measurement of the effects of the intervention.

In this chapter, various ways are described in which interventions may be randomly assigned among trial participants. The simplest method, if there are two intervention groups, is by using a procedure which is equivalent to tossing a coin to decide the allocation for each individual unit. This can either be done literally, or an equivalent procedure may be simulated using a table of random numbers or by using a computer to generate random numbers, as described in Section 2.1. In large trials, the use of such a simple randomization procedure is highly likely to ensure that there are nearly equal numbers of units allocated to the different intervention groups and the distribution of potentially confounding factors will be similar in all groups. However, if the total number of units in a study is small, such an assignment procedure may result by chance in the compositions of the different intervention groups being markedly different with respect to factors that may affect the outcome measures in the trial, or markedly unequal numbers of participants may be recruited to each intervention group. Such imbalance may arise by chance as, for example, it is possible that, if a coin is tossed ten times, it will come down heads, say, only twice. In fact, the chance that it will come down exactly heads five times and tails five times is only about 25%. For trials involving several hundreds of participants or more, any such imbalance is likely to be small and can be taken into account in the analysis of the trial. In a small trial, imbalance may make the trial more difficult to interpret, and it is advisable to design the randomization procedure to ensure balance. For this purpose, ‘restricted’ or ‘blocked’ randomization (see Section 2.2) can be used to ensure balance in group sizes. Blocked randomization also helps to achieve balance on time sequence and, in multicentre trials, study site. Stratum-matched designs (see Section 2.3) can be employed to produce balance in the composition of the groups, with respect to those variables on which the matching is based.

The techniques described in Sections 2 and 3 may be used whether the intervention is assigned to communities or to individuals. However, when communities are randomized, as in cluster randomized trials, the number of randomization units (communities) may be relatively small (often 20 or less), and more sophisticated methods of randomization have been devised to reduce sources of potential bias in the allocation of interventions in such trials. These methods are summarized in Section 3.

Whenever possible, intervention studies should be both randomized and *double-blind*, i.e. neither the participants nor the investigator should know to which group each participant has been allocated. This guards against biases that may result from knowledge of the intervention affecting the way an individual behaves, is treated, or is monitored during the trial, or assessed during, or at the end of, the trial. Blinding is discussed in Section 4. In Section 5, there is a discussion of coding systems for recording intervention allocation that may be used in trials.

2. Randomization schemes for individual participants

2.1. Unrestricted randomization

Simple random allocation of individuals between the different intervention groups is carried out most conveniently by using a computer. For example, in Microsoft Excel, the instruction ‘=RANDBETWEEN(1,3)’ will produce a random number between 1 and 3, i.e. each of the numbers 1, 2, or 3 has an equal chance of being generated. The equivalent of tossing a coin is =RANDBETWEEN(1,2). Some calculators also have a key which generates a random number on the display (usually a decimal number between 0 and 1, so that, for example, the equivalent of coin tossing would be to allocate a number less than 0.5000 as ‘heads’ and a number 0.5000 or greater as ‘tails’).

In large trials, it is common for a centralized randomization system to be used. When an investigator has decided that a participant meets the entry criteria for a trial, and the participant has given informed consent to be randomized to one of the trial interventions, the investigator telephones, or sends a text, to a central office to give the identification details for the participant, and the office then tells, or texts, the investigator to which intervention the participant has been randomly assigned or, in the case of a double-blind trial, the code for the intervention that should be administered to the participant. Systems are now commonly used whereby this process has been automated and does not require an individual to answer the telephone in the central office or for a similar automated procedure to be followed over the Internet. The advantage of this method of intervention assignment is that there is no way in which the investigator can influence the randomization procedure, and if, for example, the investigator decides not to allocate an intervention to a participant after knowing the random assignment, there is a central record of this.

For investigators who cannot set up access to a procedure for remote randomization, a frequently used alternative procedure is for a set of opaque, sealed, and numbered envelopes to be prepared, containing the intervention allocations (or possibly even the actual interventions if these are, for example, drugs). The envelopes are opened in numerical sequence, as each new person is entered into the trial. Entry criteria must be checked and eligibility satisfied before an envelope is opened, in order to exclude the possibility that the decision to accept a subject into the trial is influenced by the knowledge of the group to which he or she would be allocated. For large trials, the use of envelopes may be too cumbersome. Coding systems and alternative procedures appropriate for use in the case of ‘double-blind’ designs are discussed in Section 5.

Where the study product (for example, drug, vaccine) package is individually numbered and labelled (and randomization has been done before the numbering and labelling and where there is an indistinguishable placebo or control intervention), randomization may simply be achieved by registering each new recruit and assigning them the number on the product package.

In some circumstances, it may be better to design the randomization system, such that it is completely transparent to participants that a random allocation process is being used. A trial may be more acceptable if the trial population is involved in the randomization procedure. For example, in a trial in Ghana, the allocation of insecticide-impregnated bed-nets was randomized, such that, in some communities, all households received a bed-net immediately and, in other communities, the distribution of nets was deferred until a later time (Binka et al., 1996). At a public meeting involving all of the trial communities, the name of each community was written on a slip of paper. All the slips were put in a bucket, and a child was asked to draw some of the slips from the bucket to determine which communities received the bed-nets first. By using this procedure, it was apparent that the allocation was random and that no favouritism was operating. The fairness of the procedure was demonstrated to the population by the fact that, by

chance, the community in which the area chief resided was not selected for early bed-net allocation (much to the surprise of the population)! (Fred Binka, personal communication.)

Unrestricted randomization is often employed in large trials, as it is likely that any imbalance between the intervention groups with respect to risk factors for the occurrence of the outcomes of interest will tend to even out. Furthermore, it is possible to adjust for any residual imbalance during the analysis of the study without important loss of statistical power.

2.2. Restricted randomization

Although an unrestricted randomization procedure should lead to approximately equal numbers of participants in each group, this is not guaranteed. For example, there is more than a 5% chance that, if 20 participants are allocated to one of two groups at random, six or fewer may be allocated to one group, and 14 or more to the other. A better balance is achieved by using a 'restricted randomization' procedure, also called 'blocked randomization' or 'randomization with balance'. This procedure ensures equal numbers in each group, after there have been a fixed number of allocations. For example, the allocation procedure might be designed in blocks of ten, such that, in every ten allocations, five are to one group and five to the other. The total number of intervention groups must be a multiple of the size of the blocks.

In order to minimize the possibility that an allocation can be deduced from previous allocations, the block size should not be too small (in particular, it should not be two!), and, if possible, it should not be known to the investigator responsible for the administration of the interventions. Indeed, as far as possible, those giving the interventions should not be aware that blocking has been carried out, or, if the block size is a fixed number, the person giving the intervention would know in advance what the intervention allocation of the last individual or group in the block would be. Another safeguard is to use several different block sizes for allocating interventions in a trial. For example, in a trial with two arms, the block size might be varied, at random, between eight, ten, and 12.

Two different procedures for carrying out restricted randomization are described in Sections 2.2.1 and 2.2.2, one appropriate for small block sizes and the other appropriate for larger block sizes, say eight or more.

2.2.1. Small block sizes

If two interventions, say A and B, are to be allocated using a block size of, say four, it is possible to list all the different possible combinations of the allocations that will yield two As and two Bs. This is illustrated in Table 11.1. A number is allocated to each combination, and a random number is chosen to select a particular allocation.

The selection of each random number (between 1 and 6) generates four intervention allocations. Thus, if the random numbers 4, 5, and 1 are generated, these yield a list of twelve intervention allocations (to be assigned to participants in sequence) (Table 11.2).

2.2.2. Larger block sizes

Listing all possible combinations of allocations within a block becomes unmanageable, as the block size increases. For example, with a block size of ten, there are 252 different possible combinations, each yielding five participants in each of two intervention groups A and B. An alternative approach is necessary therefore. Suppose the block size is to be 12 and six allocations are to be made to group A and six to group B. Random numbers between 1 and 12 are generated, until six different numbers in that range have been generated (numbers that duplicate a previous one are ignored). Algorithms are easily available on the Internet to generate such random numbers. (For example, at <http://www.random.org/integers>, it is straightforward to generate X random integers between Y and Z where the user inserts values for X, Y, and Z.) Thus, we might request six random numbers between 1 and 12 and obtain 1, 2, 4, 7, 11, and 12. Then, the first, second, fourth, seventh, eleventh, and twelfth participants within the block are allocated to one of the interventions, say A, and the other participants to B. The complete sequence for the block of 12 is shown in Table 11.3.

A similar procedure, with a different set of random numbers, is used to allocate interventions in the next block (i.e. 13 to 24), and so on.

In general, it is better to choose block sizes which are not too large, in order to reduce the risk of a long sequence of individuals being allocated to the same intervention. A maximum block size of 12 is suggested.

2.3. Stratified randomization

If different subgroups of participants, say males and females, have different background rates of disease, it may be desirable to design the allocation procedure such that the interventions are equally divided in each subgroup. This may be achieved through 'stratified' randomization. The population is stratified, for example, by sex or by age group, and the allocation of the interventions is carried out separately in each stratum.

Stratification may be based on more than one factor. For example, there may be a separate allocation of interventions in each of a number of different age–sex groups. The greater the number of strata, the more complex the organization of the randomization is; in general, the number of strata should be kept small. Separate randomization lists will have to be maintained for each stratum. This may be achieved by using different sets of coloured envelopes, packages, or sticky labels for each stratum.

Stratified randomization should be considered if it is known that there are large differences in disease risk between different groups of individuals in a trial (or in response to treatment in the case of a therapeutic trial) and if it is possible to place individuals in strata corresponding to different levels of risk prior to entry to the trial. The objective of stratification is to try to include in each stratum those at similar risk of disease (or response to treatment) and to randomize between interventions separately within each stratum. In multicentre trials, randomization is often stratified on study site.

3. Randomization schemes for community or group-based interventions

As discussed in Chapter 4, trial designs have been increasingly employed in recent years, in which the unit of allocation of the intervention is a community or group, rather than an individual. These cluster randomized trials may involve the randomization of communities that can be quite large; consequently, the number of communities that can be included in a trial is often relatively small and may be of the order of 20 communities or fewer. If a method of simple unrestricted randomization is used to allocate interventions to communities, there is a reasonably high chance that there may be differences between the two groups of communities, unrelated to the interventions, that may bias the measurement of the effects of the intervention. It is common therefore to employ some method of restricted randomization in the allocation of interventions to communities (see also Chapter 4, Section 4.2).

3.1. Matched-pairs design

A matched-pairs design is a special case of stratified randomization, in which the strata are each of size two. Communities are matched into pairs, the pairs being chosen so that the two communities in a pair are as similar as possible with respect to potential confounding variables; in the absence of any intervention, the two communities would be expected to have similar incidence rates of the disease or other outcome under study. One member of each pair is assigned at random to one intervention group and one to the other. Similar matching procedures can be employed when there are more than two intervention groups. For example, with three groups, matched triplets would be employed.

Recent research on the design of cluster randomized trials has indicated that, although matched-pairs randomization remains a valid study design, other methods of randomization, such as stratified randomization or constrained (restricted) randomization, discussed in Sections 3.2 and 3.3, may generally be more appropriate design strategies (Hayes and Moulton, 2009). The major reason for this is because, if a trial is designed as a matched-pairs study, then it must be analysed as such. In technical terms, pairing reduces the number of 'degrees of freedom' that are available in the statistical comparison of the outcome measures in the intervention and comparison communities, compared to an unmatched design. This has little consequence if the number of communities is large, but, if the number is small, as is typically the case, then matching reduces the statistical power of a trial to detect an intervention effect of a given size (unless the matching factors are very closely correlated with the outcome).

3.2. Stratified design

For the reasons outlined, unrestricted randomization in a cluster randomized trial may lead to imbalance with respect to potential confounding factors between the different comparison arms of the trial, unless the number of clusters is very large. Pair matching of communities is one way of attempting to overcome this problem to ensure better balance between the arms of the trial, but this strategy may be associated with a substantial loss of statistical power. An intermediate alternative is to adopt a *stratified*, rather than a matched-pairs, design. A stratified design involves the grouping of communities into a number of strata, based on the expected rate of disease in the absence of the

intervention. For example, in a study on malaria, communities with high transmission intensity would be put into the same stratum, and those with low transmission intensity would be put into a different stratum. The communities within each stratum are then randomly allocated between the different intervention arms of the trial.

In practice, it is often challenging to decide which communities should go into the same stratum. If there are baseline rates available for the disease under study from surveillance or from a previous study, then these may provide a reasonable guide as to the expected rates in the different communities in the absence of the interventions. However, the rates of some diseases may vary substantially from year to year, and what happened in the past may not be a very good guide for what will happen in the future. Quite commonly, such rates are not available, and the investigator has the alternative of conducting a pre-trial study to estimate disease rates in each community or, based on ecological and epidemiological considerations, of making some estimate of what the rates might be. The first of these options adds to the cost of the study, whereas there may be considerable uncertainties regarding the utility and accuracy of the second approach. A fuller discussion of these issues is given in (Hayes and Moulton 2009).

A stratified design is associated with less loss of statistical power than a matched-pairs design and will assist in making the communities in the different arms of the trial more comparable with respect to potential confounding factors. There may still remain some imbalance with respect to these factors, but it is possible to adjust for this in the analysis of the trial, provided, of course, the relevant confounding factors have been measured. Methods for the analysis of cluster randomized trials and the adjustment for confounding factors are beyond the scope of this book and will generally require the input of a specialist statistician.

(Hayes and Moulton 2009) suggest that, in practical situations, it is likely that the use of three or four strata will provide most of the advantages provided by pair matching, such that communities can be very accurately paired with respect to expected disease rates during the trial. With respect to the choice of the number of strata, these authors suggest that there should be no more than two strata if there are six or fewer clusters per arm, and no more than three strata if there are 7–10 clusters per arm.

3.3. Constrained randomization design

A further method of controlling for confounding is to adopt a method known as *constrained* or *restricted* randomization. Consider a trial to be conducted in 12 communities, six of which will be allocated to the intervention under test, the remaining six serving as control communities. Using a simple unrestricted randomization design, six communities would be selected at random to receive the intervention, and the other six would serve as controls. By chance, it might happen that the six intervention communities all turn out to be close to a major highway, and the six control communities are all more distant from the highway. If the disease we are studying might be related to proximity to the highway (for example, HIV infection rates show this characteristic in some situations), then we may be rather unhappy with this particular selection of intervention communities, as there would be a priori reasons for believing there would be differences in disease rates, irrespective of the effect of the intervention we wanted to test. In these circumstances, we might reject the initial random selection of communities and select another set of random numbers to determine which our intervention communities are. While this strategy may not seem unreasonable, it is clearly dangerous to allow an investigator to override a randomization procedure if he or she does not like the result!

Constrained randomization designs aim to exclude from consideration random allocations that result in unsatisfactory imbalance between communities in the intervention and control arms. In the study already outlined, involving 12 communities, there are 924 possible different allocations of which communities comprise the six in which the intervention will be applied. Conceptually, we could imagine examining each of these possible allocations and deciding which of them we would be happy with and which would cause us concern. Suppose there were, for example, 400 for which there seemed to be a reasonable balance of confounding factors between the putative intervention and control communities. We could restrict our consideration of possible allocations to these 400, and choose one of these at random to be the one that was actually used in the trial. This is the basic principle of the constrained or restricted randomization design.

Examining all 924 possible allocations would be a considerable undertaking and would be even more difficult if the total number of communities was more than 12. It is therefore necessary to seek some more automated method of deciding which randomizations are acceptable. In practice, what is done is to define some key variables for which we wish to achieve reasonable balance across the intervention and control arms. These key variables are then compared in each of the possible randomizations, and a rule is set up to exclude a randomization if the difference between the key variables in putative intervention and control arms is more than some specified amount. Thus, the selection of

'acceptable' randomizations can be programmed into a computer, so that the selection is done automatically once the acceptability criteria for balance between the intervention and control communities have been defined.

The procedure described as a modification of simple unrestricted randomization can also be incorporated into a stratified design, so that there is a selection of acceptable possible randomizations within each stratum.

Both stratification and restricted randomization can be used to achieve good balance (avoid confounding), but stratification *also* aims to reduce between-cluster (within-stratum) variation, and hence to increase power and precision.

An example of the use of restricted randomization in the design of a trial of an adolescent sexual health intervention carried out in Tanzania (Hayes et al., 2005) is given in Box 11.1.

4. Blinding

Whenever possible, neither the participants nor the investigators should know to which intervention group each participant belongs until after the end of the trial. Such 'double-blind' designs (both the investigator and the participants are blind to the knowledge of who have received each intervention) eliminate the possibility that knowing to which intervention an individual is allocated may affect the way the individual behaves, is treated, or is monitored during the trial, or the way an individual is assessed at the end of the trial. Sometimes, a double-blind trial is not possible, and a 'single-blind' design might be used, in which the investigator knows to which group a participant belongs, but the participant does not.

'Blinded' designs are especially important when those in one of the groups under comparison are given an intervention that is expected to have no effect on the outcome of interest. To maintain blindness in these circumstances, a placebo should be used, if possible, which should look and smell as similar as possible to the intervention itself (and have a similar taste if it is being given orally). Sometimes, an identical-looking placebo cannot be obtained, and, in these circumstances, the investigator and the participants should be kept blind to which treatment is the active one. While this may be the best that can be done in some trials, it is generally undesirable. Either the participants or the investigator may form a view as to which the active treatment is (possibly erroneously), and this may affect differentially the amount of other care given to the participants or the likelihood that a participant reports apparently beneficial or harmful effects. For example, there is evidence that the colour of a tablet may affect the perceived action of a drug and seems to influence the effectiveness of a drug in some situations (de Craen et al., 1996).

For some interventions, it may be possible to preserve blindness in the initial phase of a trial, but this may be more difficult later. For example, in placebo-controlled studies of ivermectin against onchocerciasis, it was found that some participants were able to guess that they had received an active drug, rather than a placebo, because of the effect of ivermectin on other helminth infections, such as *Ascaris*, through the passage of worms in their stools, whereas those receiving placebo rarely experienced this effect. In placebo-controlled trials of BCG vaccination, most of those who have received BCG develop a lasting scar, whereas those who have received placebo do not. The possible bias that this might induce in the assessment of whether or not a participant developed leprosy, following vaccination, was overcome in a trial in Uganda by covering the vaccination site with sticking plaster for all participants before each clinical examination (Brown and Stone, 1966).

For some intervention trials, in which the unit of randomization is the community, the use of a placebo is straightforward and is no different, in principle, from the situation for an individually randomized trial. This was the case, for example, in a cluster randomized trial to assess the impact of regular vitamin A supplementation on child mortality. Those in the control communities received supplementation with an inert liquid that was administered in such a way that it was indistinguishable from the administration of vitamin A (Ghana VAST Study Team, 1993). For some interventions, however, a suitable placebo may be impossible to find. What would be a suitable placebo for an improved water supply and sanitation programme in a village, for example?

5. Coding systems

In some circumstances, it may be necessary to break the intervention code for an individual. This might arise, for example, if a severe adverse event becomes manifest and the treatment for it may be influenced by knowledge of what intervention the individual received. The coding system which is used to record which individuals received which intervention should be designed, such that, if it is necessary to break the code for one individual, the blindness of the

investigator, with respect to the interventions received by other trial participants, should be preserved. For example, if one intervention is coded A and the other B, breaking the code for one individual effectively breaks the code for all participants (if the investigator knows who has received A and who has received B). The use of a single code for each intervention is generally a poor design. It is better to have a unique code for each participant and to have a separate list linking participant numbers with the intervention allocated, or to have only a very small number of participants sharing the same code number. For example, in a BCG trial in South India for tuberculosis prevention, ampoules (each containing several doses of vaccine) were packed in boxes of three. Each box held three vials containing one of two different vaccine doses or a placebo preparation. The three ampoules were randomly coded 1, 2, and 3. The vaccine received by a participant was coded in the trial records by a combination of the box number and the ampoule number (*Tuberculosis Prevention Trial Madras, 1979*). If it had been necessary to break the vaccine code for an individual, it would only have been broken for those participants who received vaccine from the same ampoule in the same box.

The randomization list should usually be prepared in advance of the trial, and the codes assigned by someone other than the PI. If the intervention is a drug or a vaccine, the manufacturer may agree to supervise the packaging and coding, but the allocation procedure should be overseen, and the code should be held during the trial by a disinterested party. Often, the code is held by the data safety and monitoring committee (see Chapter 7, Section 4). It is also worth checking, for a random sample of the drugs or vaccines, that the codes are correct and errors have not been made in the packaging.

5.1. Individual allocations

Suppose two interventions are to be allocated between 200 individuals. A good coding scheme would be to choose 100 random numbers between 1 and 200 and allocate these codes for intervention A, say, and allocate the other 100 for intervention B (there may also be some ‘blocking’ within the total group of 200, say in blocks of size ten; see Section 2.2). When an intervention is allocated to the 127th patient in the trial, they would be given the drugs in envelope number 127, and this would be noted in their trial record. A master list of the interventions corresponding to each number would be kept in a secure place by a third party not directly connected with the trial. If it were necessary to break the code for an individual patient, the third party could do this without revealing any of the other codes to the investigator. Only at the end of the trial would the list be released to the investigator for the analysis of the results of the trial.

5.2. Group allocations

If a trial involves many thousands of participants, it may be logistically too complicated to allocate a separate treatment code number to each participant, though this will depend upon the circumstances, and, in some cases, having thousands of individual codes poses no problem. An alternative approach is to use a fixed, but not too small, number of codes for the different interventions. If there are N participants in the trial and C codes for the interventions, then breaking the code for one participant would break the codes for N/C in total. For example, the coding system used for a vaccine trial in Venezuela is given in Box 11.2. In this trial, 998 different codes were used (499 for one vaccine and 499 for the other) for about 30 000 participants. Breaking the code for one individual would break it for about 30 others (*Convit et al., 1992*).

A simpler system might be required if participants had to be given the same intervention on a number of occasions. A method that was used in a trial of ivermectin against onchocerciasis in Sierra Leone was to allocate 20 codes for ivermectin or placebo treatments (A, B, C, D, and so on) (*Whitworth et al., 1991*). The drugs were taken to the field in 20 tins, with the code letters on them (ten of which contained ivermectin, and ten contained placebo tablets), and participants were allocated to one of the 20 codes at random. If a participant was allocated, say to code E, then each time they were treated, the dose was taken from tin E. About 1000 patients were included in the trial, so that breaking the code for one individual would have also broken it for $1000/20 = 50$ others. A similar system was used in a trial of a pneumococcal vaccine in The Gambia, which involved many thousands of participants, and each participant was scheduled to receive three doses of the vaccine at different times (*Cutts et al., 2005*).

With either individual or group allocations, it is helpful if the intervention codes are on removable sticky labels that can be affixed to an individual’s form, thus minimizing the likelihood of recording errors. Where possible, the coding system should be devised so that transcription errors in recording may be detected. How this was achieved in the leprosy vaccine trial in Venezuela is illustrated in Box 11.2. More commonly now, bar codes are used to identify

interventions in trials using drugs or vaccines, and, provided that suitable computer systems are set up, this should eliminate the possibility of transcription errors.

References

- Binka, F. N., Kubaje, A., Adjuik, M., et al. 1996. Impact of permethrin impregnated bednets on child mortality in Kassena-Nankana district, Ghana: a randomized controlled trial. *Tropical Medicine and International Health*, **1**, 147–54.10.1111/j.1365-3156.1996.tb00020.x [PubMed: 8665378] [CrossRef]
- Brown, J. A. and Stone, M. M. 1966. B.C.G. vaccination of children against leprosy: first results of a trial in Uganda. *BMJ*, **1**, 7–14.10.1136/bmj.1.5478.7 [PMC free article: PMC1845255] [PubMed: 5900912] [CrossRef]
- Convit, J., Sampson, C., Zuniga, M., et al. 1992. Immunoprophylactic trial with combined Mycobacterium leprae/BCG vaccine against leprosy: preliminary results. *Lancet*, **339**, 446–50.10.1016/0140-6736(92)91056-E [PubMed: 1346818] [CrossRef]
- Cutts, F. T., Zaman, S. M., Enwere, G., et al. 2005. Efficacy of nine-valent pneumococcal conjugate vaccine against pneumonia and invasive pneumococcal disease in The Gambia: randomised, double-blind, placebo-controlled trial. *Lancet*, **365**, 1139–46.10.1016/S0140-6736(05)71876-6 [PubMed: 15794968] [CrossRef]
- De Craen, A. J., Roos, P. J., Leonard De Vries, A., and Kleijnen, J. 1996. Effect of colour of drugs: systematic review of perceived effect of drugs and of their effectiveness. *BMJ*, **313**, 1624–6.10.1136/bmj.313.7072.1624 [PMC free article: PMC2359128] [PubMed: 8991013] [CrossRef]
- Ghana VAST Study Team 1993. Vitamin A supplementation in northern Ghana: effects on clinic attendances, hospital admissions, and child mortality. *Lancet*, **342**, 7–12.10.1016/0140-6736(93)91879-Q [PubMed: 8100345] [CrossRef]
- Hayes, R. J., Chagalucha, J., Ross, D. A., et al. 2005. The MEMA kwa Vijana project: design of a community randomised trial of an innovative adolescent sexual health intervention in rural Tanzania. *Contemporary Clinical Trials*, **26**, 430–42.10.1016/j.cct.2005.04.006 [PubMed: 15951245] [CrossRef]
- Hayes, R. J. and Moulton, L. H. 2009. *Cluster randomized trials*. Boca Raton, Fla.; London: Chapman & Hall/CRC.10.1201/9781584888178 [CrossRef]
- Schulz, K. F. 1995. Subverting randomization in controlled trials. *JAMA*, **274**, 1456–8. [PubMed: 7474192]
- Tuberculosis Prevention Trial Madras. 1979. Trial of BCG vaccines in south India for tuberculosis prevention. *Indian Journal of Medical Research*, **70**, 349–63. [PubMed: 535944]
- Whitworth, J. A. G., Gilbert, C. E., Mabey, D. M., Maude, G. H., Morgan, D., and Taylor, D. W. 1991. Effects of repeated doses of ivermectin on ocular onchocerciasis: community-based trial in Sierra Leone. *Lancet*, **338**, 1100–3.10.1016/0140-6736(91)91963-U [PubMed: 1682543] [CrossRef]

Tables

Table 11.1 Example of allocation rule for a block size of four, with two intervention groups A and B

Allocation	Corresponding random number
AABB	1
BBAA	2
ABAB	3
BABA	4
ABBA	5
BAAB	6

Table 11.2 Example of random allocation to two groups using a block size of four

Block number	1	2	3
Random number	4	5	1
Allocation sequence	BABA	ABBA	AABB

Table 11.3 Example of random allocation to two groups using a block size of 12

Participant	1	2	3	4	5	6	7	8	9	10	11	12
Intervention	A	A	B	A	B	B	A	B	B	B	A	A

Boxes

Box 11.1 Use of restricted randomization in a community randomized trial of an adolescent sexual health intervention in Tanzania

In this trial, carried out to evaluate the impact of a multi-component sexual health intervention on HIV and other adverse outcomes among adolescents in Tanzania, the 20 rural study communities were grouped into three strata, based on their expected risk of HIV infection (Hayes et al., 2005). There were six communities in the low-risk stratum, eight in the medium-risk stratum, and six in the high-risk stratum.

There is a total of 28 000 ways of assigning half the communities in each stratum to the intervention arm and half to the control arm. Because the total number of communities is quite small, not all of these 28 000 allocations would provide a good balance of key characteristics across treatment arms. Restricted randomization was therefore used to achieve an acceptable balance by applying the following criteria:

- ◆ mean HIV prevalence in each treatment arm within 0.075% of overall mean
- ◆ mean prevalence of *Chlamydia trachomatis* (CT) infection in each treatment arm within 0.1% of overall mean
- ◆ two of the 20 communities were close to gold mines, and one of these was to be allocated to each treatment arm
- ◆ even distribution of intervention communities across the four administrative districts in which the trial was carried out.

HIV and CT prevalence were based on an initial survey of young people carried out in each study community. Prevalences of HIV and CT (also an STI) were assumed to be correlated with sexual behaviour in the study communities and therefore to be predictors of the risk of acquiring HIV infection during the trial. HIV prevalence is often increased in mining communities, and it was important to ensure that one mining community was allocated to each treatment arm. Finally, ensuring an even distribution of intervention communities across districts helped to ensure that the trial was acceptable to local leaders.

A computer program was used to check each of the 28 000 possible allocations against the balance criteria, and 953 allocations satisfied the criteria and were listed. One of these was chosen randomly at a public randomization ceremony.

Source: data from Hayes, R. J., et al., The MEMA kwa Vijana project: design of a community randomised trial of an innovative adolescent sexual health intervention in rural Tanzania, *Contemporary Clinical Trials*, Volume 26, Issue 4, pp. 430–42, Copyright © 2005 Elsevier Inc. All rights reserved.

Box 11.2 Assignment of check letter for three-digit vaccine code

The coding system described was that used in a leprosy vaccine trial conducted in Venezuela (Convit et al, 1992). Randomization was to one of two vaccines.

The vaccine vials were labelled with a number between 1 and 998. A total of 499 of these numbers were allocated at random for one vaccine, and the other 499 for the other vaccine. A check letter was added to each number, so that transcription errors would stand a high chance of being detected. The code was devised, such that every possible permutation of the same three digits in a number had a different check letter, as illustrated:

In some countries, number 1 and number 7 are distinguished clearly when written, as it is the custom for the number 7 to have a horizontal stroke put through it. In other countries, however, this is not the custom, and there is a danger that these numbers will be confused. In such cases, it would be advisable to change the check coding system, such that, if a 1 is confused with a 7, or vice versa, the check letter will enable the error to be detected. Thus, the system outlined might be modified, as indicated:

001A	010B	100C			
002D	020E	200F			
.					
.					
009M	090N	900P			
010B—already allocated—see line 1					
011R	101S	110T			
.					
.					
123W	132X	213Y	231A	312B	321C
124D	142E	214F	241G	412H	421J
etc.					

001A	010B	100C	007D	070E	700F			
002G	020H	200J						
003K	.	.						
.								
.								
.								
011R	101S	110T	017V	071X	107Y	170A	701B	710C
077D	707E	770F						
012G	etc.							

Source: data from Peter Smith, (personal communication).

© London School of Hygiene and Tropical Medicine 2015.

This is an open access publication. Except where otherwise noted, this work is distributed under the terms of the Creative Commons Attribution NonCommercial 4.0 International licence (CC BY-NC), a copy of which is available at <http://creativecommons.org/licenses/by-nc/4.0/>. Enquiries concerning use outside the scope of the licence terms should be sent to the Rights Department, Oxford University Press, at the address above.

Monographs, or book chapters, which are outputs of Wellcome Trust funding have been made freely available as part of the Wellcome Trust's open access [policy](#)

Bookshelf ID: NBK305495