Clinical Epidemiology

A clinician's guide for conducting randomized trials in individual patients

Gordon Guyatt, MD
David Sackett, MD
Jonathan Adachi, MD
Robin Roberts, MSc
John Chong, MD
David Rosenbloom, Pharm D
Jana Keller, BSc

In determining optimal treatment for a patient conventional trials of therapy are susceptible to bias. Large-scale randomized trials can provide only a partial guide and have not been or cannot be carried out for most clinical disorders. However, randomized controlled trials (RCTs) in individual patients (N of 1 RCTs) may in some circumstances provide a solution to this dilemma. In an N of 1 RCT a patient undergoes pairs of treatment periods (one period of each pair with the active drug and one with matched placebo, assigned at random); both the patient and the clinician are blind to allocation, and treatment targets are monitored. N of 1 RCTs are useful for chronic, stable conditions for which the proposed treatment, which has a rapid onset of action and ceases to act soon after it is discontinued, has shown promise in an open trial of therapy. The monitoring of treatment targets usually includes quantitative measurement of the patient's symptoms with the use of simple patient diaries or questionnaires. Pairs of treatment periods are continued until effectiveness is proved or refuted. The cooperation of a pharmacy is required for the preparation of matching placebos and conduct of the trial. Formal statistical analysis may be helpful for interpreting the results. The practical approach presented in this paper allows clinicians to conduct their own N of 1 RCTs.

Dans la recherche du meilleur traitement les essais classiques ne sont pas toujours exempts de biais. Ceux qui ont porté sur de nombreux sujets, même choisis au hasard, ne donnent pas toutes les réponses pour la plupart des maladies ou bien ils n'ont pas été faits ou bien ils sont impossibles. La difficulté sera parfois tournée par un essai comparatif sur un seul sujet agissant comme son propre témoin (ECT [N = 1]) si on le traite plusieurs fois, alternativement et à double insu, par le médicament actif et par un placébo, les périodes pendant lesquelles on donne l'un et l'autre étant déterminées au hasard, et qu'on est à l'affût des effets thérapeutiques visés. Ce genre d'essai convient, dans une maladie chronique et stabilisée, à un traitement dont l'effet se manifeste rapidement quand on l'institue et cesse peu après qu'on ne l'administre plus et qui a paru prometteur lors d'essais préliminaires. La constatation des effets thérapeutiques comporte habituellement la mesure quantitative des symptômes, soit par un journal que tient le malade, soit par un questionnaire. L'alternance des périodes de traitement au médi-

cament ou au placébo se poursuit jusqu'à preuve de l'efficacité ou de l'inefficacité du médicament en question. La préparation des placébos et la logistique de l'essai exigent la collaboration du pharmacie. L'interprétation des résultats se trouve parfois facilitée par l'analyse statistique. De la façon que nous avons décrite, le médecin est à même de faire ses propres ECT (N = 1).

When deciding on optimal treatment for a patient, clinicians often cannot rely on the results of randomized controlled trials (RCTs). The relevant trial may never have been done, or the results of those that have been done may not apply to that patient.

Under these circumstances clinicians typically conduct the time-honoured "trial of therapy", in which the patient is given a treatment, and the subsequent clinical course determines whether the treatment is judged effective and continued. Many factors may mislead physicians conducting conventional therapeutic trials: the placebo effect, the natural history of the illness, the expectations of the patient and the clinician about the treatment effect, and the desire of the patient and the clinician not to disappoint one another.[1]

These pitfalls can be avoided in trials of therapy by safeguards that permit the natural, untreated course of the disorder to be observed and keep both the patient and his or her clinician "blind" to when active treatment is being administered. Such safeguards are routine in large-scale RCTs involving dozens or hundreds of patients. We describe our approach to transferring these safeguards to the evaluation of therapy in the individual patient.

We have previously detailed the rationale and general approach to an N of 1 RCT.[1] Although such "single-subject experiments" are an established research method in psychology,[2-4] they have only recently been applied in medical practice.[1,5] Our experience in over 50 N of 1 RCTs, coupled with the number of inquiries we have received about their use, has convinced us that such trials may be widely applicable in clinical practice. Accordingly, in this paper we go beyond our previous description of the method and provide sufficient detail to allow clinicians to plan and execute their own N of 1 RCTs.

Although there are many ways of conducting N of 1 RCTs the method we have found most widely applicable can be summarized as follows.

• A clinician and a patient agree to test a therapy (the "experimental therapy") for its ability to reduce or control the symptoms, signs or other manifestations (the "treatment targets") of the patient's ailment.

• The patient then undergoes pairs of treatment periods; one period of each pair applies the experimental therapy, and the other applies either an alternative treatment or a placebo. The order of the two periods within each pair is randomized by a coin toss or another method that ensures that each period has an equal chance of applying the experimental or the alternative therapy.

• Whenever possible both the clinician and the patient are blind to the treatment being given during either period.

• The treatment targets are monitored (often through the use of a patient diary) to document the effect of the treatment being applied.

• Pairs of treatment periods are replicated until the clinician and the patient are convinced that the experimental therapy is effective, is harmful or has no effect on the treatment targets.

To help clinicians implement an N of 1 RCT we will address each step with a question that must be answered before they proceed to the next step (Table I).

Guidelines for performing an N of 1 RCT

Is an N of 1 RCT indicated for this patient?

Because N of 1 RCTs are unnecessary for some ailments (such as self-limited illnesses) and unsuited for some treatments (such as operations) it is important to determine at the outset whether an N of 1 RCT is really indicated for the patient and treatment in question. If it is appropriate the answers to each of the following questions should be Yes.

Is the effectiveness of the treatment really in doubt?

One or several RCTs may have shown that the treatment is highly effective. If one is unsure whether such trials have been undertaken, efficient strategies for searching the medical literature are available.[6-11] However, if a substantial proportion of the subjects in such trials have proved un-

---

Table I — Guidelines for performing an N of 1 randomized controlled trial (RCT)

Is an N of 1 RCT indicated for this patient?
  Is the effectiveness of the treatment really in doubt?
  Will the treatment, if effective, be long-term?
  Is the patient eager to collaborate in designing and carrying out an N of 1 RCT?
Is an N of 1 RCT feasible in this patient?
  Does the treatment have a rapid onset?
  Does the treatment stop acting soon after it is discontinued?
  Is an optimal duration of treatment feasible?
  Can clinically relevant targets be measured?
  Can sensible criteria for stopping the trial be established?
  Is an unblinded run-in period necessary?
Is the trial feasible in my practice setting?
  Is there a pharmacist who can help me?
  Are strategies for interpreting the data in place?
Is the trial ethical?

responsive an N of 1 RCT may still be appropriate.

Alternatively, a patient may have exhibited such a dramatic response to the treatment that both the clinician and the patient are convinced that it works. N of 1 trials really aren't necessary in such cases and are best reserved for the following situations.

- When neither the clinician nor the patient is confident that a treatment is really providing benefit.
- When the clinician is uncertain whether a treatment that hasn't yet been started will work in a particular patient.
- When the patient insists on taking a treatment that the clinician thinks is useless or potentially harmful (and mere words won't change the patient's mind).
- When a patient is experiencing symptoms that the clinician and the patient suspect are medication side effects, but neither is certain.
- When neither the clinician nor the patient is confident of the optimal dose of a medication or replacement therapy (such as thyroxine for a patient with hypothyroidism).

Will the treatment, if effective, be long-term?

If the underlying condition is self-limited and treatment will only be short-term an N of 1 RCT may not be worth while. N of 1 RCTs are most useful for chronic conditions for which maintenance therapy is likely to be continued for long periods.

Is the patient eager to collaborate in designing and carrying out an N of 1 RCT?

An N of 1 RCT is indicated only when patients can fully understand the experiment and are enthusiastic about participating. By its nature the N of 1 RCT is a cooperative venture between clinician and patient.

Is an N of 1 RCT feasible in this patient?

The clinician may wish to determine treatment effectiveness in an individual patient, but the ailment or the treatment may not lend itself to the N of 1 approach. Once again, for the N of 1 approach to be feasible the answer to each of the following questions must be Yes.

Does the treatment have a rapid onset?

An N of 1 RCT is much easier to do when a treatment, if effective, begins to act within hours. Although it may be possible to do a trial with drugs of longer latency (such as gold or penicillamine for rheumatoid arthritis or tricyclics for de-

pression) the requirement for very long treatment periods may become prohibitive.

Does the treatment stop acting soon after it is discontinued?

Treatments whose effects abruptly cease when they are withdrawn are the most suitable for N of 1 RCTs. If the treatment continues to act long after it is stopped, a prolonged washout period may be necessary. If the washout period is more than a few days the feasibility of the trial is compromised. Similarly, treatments that have the potential to "cure" the underlying condition or at least lead to a permanent change in the treatment target are not suitable for N of 1 RCTs. This is true, for example, of most behavioural interventions and treatments in such areas as physiotherapy and occupational therapy. The N of 1 trial may be modified in these situations, but such modification is beyond the scope of this discussion.

Is an optimal duration of treatment feasible?

Although short treatment periods boost the feasibility of an N of 1 RCT, the periods may have to be long to be valid. For example, if active therapy takes a few days to reach its full effect and a few days to stop acting once it is discontinued, treatment periods of sufficient duration to avoid distortion from these delayed peak effects and washout periods are required. Thus, in our N of 1 RCTs of theophylline in asthma the treatment periods are at least 10 days — 3 days to allow the drug to equilibrate or wash out and 7 days to monitor the patient's response to treatment.

Since many N of 1 RCTs will test a treatment's ability to prevent or blunt attacks or exacerbations of disease (such as migraines or seizures) each treatment period must be long enough to include an attack or exacerbation if one is going to occur. A rough rule of thumb, the "inverse rule of 3", tells us that if an event occurs, on average, once every x days, we need 3x days to be 95% confident of seeing at least one event. Thus, in a patient with familial Mediterranean fever who has an attack, on average, once every 2 weeks the treatment periods need to be 6 weeks long.

The clinician may not want the patient to take responsibility for crossing over from one treatment period to the next or may need to examine the patient at the end of each period. Thus, the clinician's office schedule and the patient's travel considerations may dictate the length of each treatment period.

Can clinically relevant treatment targets be measured?

The treatment targets, or outcome measures,

usually go beyond a set of physical signs (e.g., the rigidity and tremor of parkinsonism and the jugular venous distension and pulmonary crackles of congestive heart failure), a laboratory test (e.g., measurement of the blood glucose level or the erythrocyte sedimentation rate) or a measure of patient performance (e.g., recordings of the peak airflow or the score on a 6-minute walk test), each of which is only an indirect measure of the patient's prognosis and quality of life.

In most situations it is not only possible but preferable to measure a patient's symptoms, well-being or quality of life directly. The principles of quality-of-life measurement can be applied simply to N of 1 RCTs.[12-14] One asks the patient to identify his or her most troubling symptoms or problems, and then one decides which of the symptoms or problems are likely to respond to the experimental treatment. This responsive subset of symptoms or problems forms the basis of a self-administered patient diary or questionnaire.

For example, a patient with chronic airflow limitation identified his problem as the shortness of breath he experiences while walking up stairs, bending or vacuuming. A patient with fibrositis (to whom we shall return later) identified fatigue, aches and pains, morning stiffness and sleep disturbance as the treatment targets for her illness.

The questionnaire can be presented in several formats. A daily diary is best for some; a weekly summary may be better for others. For some targets patients can quantify their symptoms with a visual analogue scale — that is, a straight line, the ends of which present the extremes of the target being measured (e.g., no shortness of breath and extreme shortness of breath). For other targets and patients graded descriptions from no symptoms to severe symptoms (e.g., no shortness of breath, a little shortness of breath, moderate shortness of breath and extreme shortness of breath) are sometimes easier. Constructing simple symptom questionnaires is not difficult, and it allows the patient and the clinician to collaborate in the quantification of symptoms upon which the analysis of the N of 1 RCT often relies.

Whatever formats are chosen for measuring treatment targets, our experience has taught us that patients should rate each of the targets at least twice during each study period. The identifying patient information and the ratings for the treatment targets can often be combined onto a one-page form; Fig. 1 displays a form for an N of 1 RCT that examines the effectiveness of a new drug, ketanserin, in treating Raynaud's phenomenon. More detailed guides for constructing the sort of simple questionnaires required for N of 1 RCTs are available.[15,16]

The measurement of a patient's symptoms may also include the side effects of treatment. A patient diary or questionnaire can be used to measure nausea, gastrointestinal disturbances, dizziness and other common side effects. This is particularly important when a side effect may lead to "unblinding" of the trial. In N of 1 RCTs designed to determine whether medication side effects are responsible for a patient's symptoms (e.g., whether a patient's fatigue is caused by the antihypertensive agent he is taking) side effects become the primary treatment targets.

Can sensible criteria for stopping the trial be established?

One of the advantages to not specifying the number of pairs of treatment periods in advance is that the clinician and the patient can stop the trial any time. For example, if there is a dramatic difference in the treatment target between the two periods of the first pair the clinician and patient may want to stop the trial immediately. Alterna-

N of 1 randomized controlled trial — data sheet 1

Physician:_____
Patient: _____
Sex: 1) Male 2) Female Date of birth ___ ___ ___
Diagnosis: _____
Occupation: _____
Present medications: _____
Trial medication: ketanserin   Dose: _____
Duration of study periods: 2 weeks
Outcomes (symptom ratings): _____
Informed consent obtained (please sign)_____
Answers to symptom questions, pair 1, period 1:
1. How many episodes of Raynaud's phenomenon did you have in the last week?
 First week (to be completed on ___ ___) ___
 Second week (to be completed on ___ ___) ___
2. On average, in comparison with your usual episodes, how long were the attacks?
 1. Very long; as long as or longer than they have ever been
 2. Very long; almost as long as they have ever been
 3. Longer than usual
 4. About as long as usual
 5. Not as long as usual
 6. Not nearly as long as usual
 7. Very short; as brief as or briefer than they have ever been
Write in the best number for each week.
 First week (to be completed on ___ ___) ___
 Second week (to be completed on ___ ___) ___
3. On average, in comparison with your usual episodes, how severe were the attacks?
 1. Very bad; as severe as or more severe than they have ever been
 2. Very bad; almost as severe as they have ever been
 3. More severe than usual
 4. About as severe as usual
 5. Not as severe as usual
 6. Not nearly as severe as usual
 7. Very mild; as mild as or milder than they have ever been
Write in the best number for each week.
 First week (to be completed on ___ ___) ___
 Second week (to be completed on ___ ___) ___

Fig. 1 — Data collection form for an N of 1 randomized controlled trial

tively, if there is a minimal difference between the two periods of each pair the clinician and patient may need three, four or even five or more pairs before confidently concluding that the treatment is or is not effective.

If, on the other hand, one wishes to conduct a formal statistical analysis of data from the N of 1 RCT the analysis will be strengthened considerably if the number of pairs is specified in advance. (This issue is discussed further in the section on interpreting the results of N of 1 RCTs.)

Whether or not one specifies the number of treatment periods in advance it is advisable to have at least two pairs of treatment periods before breaking the trial. Too many conclusions drawn after a single pair will be either false positive (that the treatment is effective when it isn't) or false negative (that the treatment is not effective when it is). Indeed, we recommend that clinicians resist the temptation to break the code until they are quite certain that they are ready to terminate the study.

Is an unblinded run-in period necessary?

A preliminary run-in period with active therapy, during which both the clinician and the patient know that active therapy is being received, could save a lot of time. After all, if there is no hint of a response during such an open trial or if intolerable side effects occur an N of 1 RCT may be fruitless or impossible. For example, we prepared for a double-blind N of 1 RCT of methylphenidate hydrochloride in a child with hyperactivity only to find a dramatic increase in agitation over the first 2 days of the first study period, during which the patient was receiving the active drug; thus, the study was abruptly terminated. An open run-in period may also be used to determine the optimal dose of the medication.

Is the trial feasible in my practice setting?

Clinicians may answer Yes to the preceding questions and still be unsure about how to proceed. In the following section we describe the mechanisms that will ensure the feasibility of an N of 1 RCT in a given practice.

Is there a pharmacist who can help me?

An N of 1 RCT that incorporates all the aforementioned safeguards against bias and misinterpretation requires collaboration between the clinician and a pharmacist or pharmacy service. Preparation of placebos identical in appearance, taste and texture to the active medication is required. Occasionally pharmaceutical firms can supply such placebos. More often, however, clinicians will want their local pharmacist to repackage the active medication; for example, if it comes as a pill

it can be crushed and repackaged in a capsule. Identical placebo capsules can be filled with lactose. While somewhat time-consuming, the preparation of placebos is not technically difficult. Our average cost for preparing medication for N of 1 RCTs in which placebos have not been available from the pharmaceutical company has been $125, which has been paid out of research funds. It could be argued that as N of 1 RCTs become part of conventional clinical practice, such costs should be picked up by a patient's drug benefit plan. We are exploring this possibility.

The pharmacist is also charged with preparing the randomization schedule (which requires nothing more than a coin toss for each pair of treatment periods) so that the clinician and the patient remain blind to allocation.

The pharmacist can provide information on the time to onset of action and the washout period and therefore help with decisions about the duration of the study periods. The pharmacist can also help monitor compliance and drug absorption. Both pill counts and the measurement of serum drug levels at the end of each treatment period can help establish that the patient is conscientiously taking the medication throughout the trial.

Are strategies for interpreting the trial data in place?

Once the data on the treatment targets are carefully gathered how will they be interpreted? One approach is to simply plot the data and examine the results visually. Such evaluation has a long and distinguished record in the psychology literature on single-subject designs and is strongly advocated by some practitioners of single-subject studies.[2-4] Visual inspection is simple and easy. Its major disadvantage is that it is open to bias.

An alternative approach is to use a statistical test of significance. The simplest test, the sign test,[17] is based on the likelihood that a patient will prefer active treatment in each pair of treatment periods. This situation is analogous to the likelihood that heads will come up repeatedly in a series of coin tosses. For example, the likelihood that a patient will prefer active to placebo treatment during three consecutive pairs of treatment periods if the treatment was ineffective would be $1/2 \times 1/2 \times 1/2 = 1/8$, or 0.125. The disadvantage of this approach is that it lacks power: five pairs of treatment periods are necessary before there is any chance that conventional levels of statistical significance will be reached.

A second statistical strategy is the Student's t-test. The t-test offers increased power because not only the direction but also the strength of the treatment effect in each pair are taken into account. The disadvantage of this test is that it makes additional assumptions about the data that may not be valid. The assumption of greatest concern is that observations are independent of one another;

that is, a patient is equally likely to feel good or bad on a particular day irrespective of whether he or she felt good or bad the day before. While some autocorrelation (i.e., the data are not independent) is likely in many N of 1 RCTs, its impact can be reduced if one uses the average of all measurements in a given period, rather than the individual measurements, in the statistical analysis. The paired design of the N of 1 RCT will further reduce the impact of any autocorrelation.

If a statistical test is used to interpret the data there is another potential problem. If the clinician and patient use the results from the study to determine when to stop the trial, the true p value may be inflated above the nominal p value. Therefore, the number of periods should be determined before the study begins.

For a paired *t*-test a single score is derived for each pair of treatment periods by subtracting the mean score of the placebo period from the mean score for the active treatment period. The differences constitute the data for the paired *t*-test; the degrees of freedom are simply the number of pairs minus 1. Statistical packages are available for any programmable pocket calculator or microcomputer that will allow calculation of the p value within seconds.

An example of the results of an N of 1 RCT is presented in Table II. This trial analysed the effectiveness of amitriptyline hydrochloride, 10 mg given at bedtime, in a patient with fibrositis. Each week the patient rated the severity of each of several symptoms, including fatigue, aches and pains, and sleep disturbance, on a seven-point scale (a higher score represented better function). Each treatment period lasted 4 weeks, and three pairs were undertaken.

The first step in analysing the results of the study is to calculate the mean score for each period (presented in the last column of Table II). In each pair the score favoured the active treatment. The sign test indicates that the probability that this result would occur by chance if the treatment was ineffective is $\frac{1}{2} \times \frac{1}{2} \times \frac{1}{2} = \frac{1}{8}$, or 0.125.

However, this analysis ignores the magnitude and consistency of the difference between active and placebo treatment. A paired *t*-test takes these factors into account. All that is required is to punch into a calculator or microcomputer the pairs of results: 4.68 and 4.22, 5.01 and 4.07, and 5.04 and 4.18. The *t* value is 5.07, and there are two degrees of freedom; the associated p value is 0.037. This analysis makes us considerably more confident that the consistent difference in favour of active treatment is unlikely to have occurred by chance.

Clinicians and statisticians may remain uncomfortable with our suggested approach to analysing the data from N of 1 RCTs. The use of N of 1 RCTs to improve patient care does not depend on the statistical analysis of the results. Even if statistical analysis is not used, the strategies of randomization, double-blinding, replication and quantitation of outcomes, when accompanied by careful visual inspection of the data, will still allow a much more rigorous assessment of treatment effectiveness than is possible in conventional clinical practice.

Is the trial ethical?

Is an N of 1 RCT a clinical or a research undertaking? If the former, is it the sort of clinical procedure, analogous to an invasive diagnostic test, that requires written informed consent? We think that an N of 1 RCT can and should be a part of routine clinical practice. But, like all medical innovations, it may require a period of experimentation and study before it is accepted by clinicians and ethics committees.

Nevertheless, several ethical issues are important. We believe that patients should be fully informed of the nature of the study in which they are participating and that there should be no element of deception in the use of placebos as part of the study. Written informed consent should be obtained (the form we use is available on request). Patients should be aware that they can terminate the trial at any time without jeopardizing their care or their relationship with their physician. Finally, follow-up should be close enough to prevent any deleterious consequences of use or withdrawal of therapy.

Discussion

Clinicians can incorporate some of the principles of N of 1 RCTs into their practices without adopting the full rigour of the approach presented here. Medication can be repeatedly withdrawn and reintroduced in an open or unblinded fashion. Symptoms and physical findings can be carefully quantitated. However, without the additional feature of double-blinding, both the placebo effect and the clinician's and patient's expectations can bias the results.

Table II — Results of an N of 1 RCT in a patient with fibrositis

Treatment*	Treatment period; severity score†				Mean score
	Week 1	Week 2	Week 3	Week 4	
Pair 1					
Active	4.43	4.86	4.71	4.71	4.68
Placebo	4.43	4.00	4.14	4.29	4.22
Pair 2					
Active	4.57	4.89	5.29	5.29	5.01
Placebo	3.86	4.00	4.29	4.14	4.07
Pair 3					
Active	4.29	5.00	5.43	5.43	5.04
Placebo	3.71	4.14	4.43	4.43	4.18

*The active drug was amitriptyline hydrochloride.
†Higher scores represent better function.

N of 1 RCTs are feasible as part of routine clinical practice. They do, however, require that additional time and energy be devoted to the patient and to organizing the study. Busy clinicians may find that the time requirement prohibits the conduct of N of 1 trials.

To help clinicians in our area we have established an "N of 1 service", which operates in two ways. For those who would like to be directly involved in an N of 1 RCT but who do not have the time the N of 1 service determines the study design, prepares the questionnaire and other forms required, and undertakes the formal analysis. The clinician then carries out the N of 1 RCT, reviewing the patient's status at the end of each period or pair of periods. Alternatively, the clinician may simply refer the patient to the N of 1 service, specifying the disorder, the treatment regimen and the possible treatment targets. The N of 1 service then carries out the trial and returns the patient and the results to the referring clinician.

N of 1 RCTs are being used in our region to determine medication requirements in particularly challenging situations (such as in a young woman with familial Mediterranean fever in whom colchicine was being given with questionable effect or in a man with an apparent remission of Menière's disease after treatment with phenytoin had been started) or in the investigation of experimental treatments (e.g., trimebutine maleate in irritable bowel syndrome). In addition, one of us (G.G.) is routinely using the method to determine the need for inhaled and oral bronchodilators in patients with chronic airflow limitation and asthma, and another of us (J.A.) is conducting them before consigning patients with fibrositis to long-term use of amitriptyline.

The N of 1 RCT clearly has great potential for improving both the quality of medical care and the judicious use of expensive and potentially harmful medication in patients with chronic disease. We believe that with the use of the guidelines offered here clinicians will find the conduct of N of 1 RCTs feasible, highly informative and fun.

References

1. Guyatt GH, Sackett D, Taylor DW et al: Determining optimal therapy — randomized trials in individual patients. *N Engl J Med* 1986; 314: 889–892

2. Kratchwill TR (ed): *Single Subject Research: Strategies for Evaluating Change*, Acad Pr, Orlando, Fla, 1978: 316

3. Barlow DH, Hersen M: *Single Case Experimental Designs: Strategies for Studying Behavior Change*, 2nd ed, Pergamon, New York, 1984: 419

4. Kazdin AE: *Single-Case Research Designs: Methods for Clinical and Applied Settings*, Oxford U Pr, New York, 1982: 368

5. McLeod RS, Taylor DW, Cohen Z et al: Single-patient randomised clinical trial: use in determining optimum treatment for patient with inflammation of Kock continent ileostomy reservoir. *Lancet* 1986; 1: 726–728

6. Haynes RB, McKibbon KA, Fitgerald D et al: How to keep up with the medical literature. I. Why try to keep up and how to get started. *Ann Intern Med* 1986; 105: 149–153

7. Idem: How to keep up with the medical literature. II. Deciding which journals to read regularly. Ibid: 309–312

8. Idem: How to keep up with the medical literature. III. Expanding the volume of literature that you read regularly. Ibid: 474–478

9. Idem: How to keep up with the medical literature. IV. Using the literature to solve clinical problems. Ibid: 636–640

10. Idem: How to keep up with the medical literature. V. Personal computer access to the medical literature. Ibid: 810–824

11. McKibbon KA, Haynes RB, Fitzgerald D et al: How to keep up with the medical literature. VI. How to store and retrieve articles worth keeping. Ibid: 978–984

12. Kirshner B, Guyatt GH: A methodological framework for assessing health indices. *J Chronic Dis* 1985; 38: 27–36

13. Guyatt GH, Bombardier C, Tugwell PX: Measuring disease-specific quality of life in clinical trials. *Can Med Assoc J* 1986; 134: 889–895

14. Guyatt GH, Berman LB, Townsend M et al: Should study subjects see their previous responses? *J Chronic Dis* 1985; 38: 1003–1007

15. Woodward CA, Chambers LW: *Guide to Questionnaire Construction and Question Writing*, Can Public Health Assoc, Ottawa, 1983

16. Sudman S, Bradburn NM: *Asking Questions: a Practical Guide to Questionnaire Design*, Jossey-Bass, San Francisco, 1982

17. Conover WJ: *Practical Nonparametric Statistics*, Wiley, New York, 1971: 121–126

A curious paradox

We are young when we expect variety, and indeed anything that promises variety or seeks change has youth. It is a curious paradox that we desire stability for our plans and require change for our souls' sake.

— *Alan Gregg (1890–1957)*