The Priority Search ranking process

Much has been written about the use of *sequential paired comparison* to establish a rank order. It was originally designed to assist people who find making choices difficult, - perhaps because of brain damage, psychosis or low ability - to choose between a small number of alternatives. The method involves presenting all possible pairs of alternatives and asking the respondent to make a dichotomous choice between each pair. Whilst it is then relatively simple to construct a rank order from this, the number of pairs increases with the square of the number of items, becoming rapidly unusable with even modest numbers of items. For example, ten items require 45 pairings, 20 items require 190, and 30 items require 435.

The Priority Search process allows respondents to compare each pair not dichotomously but using a scale. The addition of this scale gives more information per pair, and as a result the number of pairings needed is reduced considerably:

> A uniquely ranked list of n items comprises $\log_2(n!)$ bits of information.
>
> A set of 3 pairings per item on a scale of P points comprises $\log_2(P^{1.5n})$ bits, and for even small values of P the value of $P^{1.5n}$ excedes $n!$ over a usable range of items.

In order to extract a rank order from the resulting partial set of all possible pairings it is necessary to be able to relate each item to all others. Consider a set of ten items paired as follows:

A - F B - G C - H D - I E - J	In this case, we know how A relates to F, B to G, etc, but we have no information about how A relates to any item other than F, or B to any item other than G, etc.

A - F B - G C - H D - I E - J B - F C - G D - H E - I A - J	By creating a second set of pairings with the left hand column frame shifted, a chain results: At the top, A is compared with F, which below is compared with B; B is compared with G, which is compared with C, and so on. In this way the position of any item relative to any other can be determined.

Such a design is known as a *reduced subset cyclic design*. Two sets of pairings arranged as above will allow a perfect rank order to be calculated if the input to the system comprises mathematically precise data. The Priority Search process adds a third, different set of pairings; this allows more information to be extracted and this is sufficient to cope with the imprecision which is inherent in subjective ratings.

The use of paired comparisons

We are frequently asked questions about the how the paired comparison system works. Sometimes, users of the Priority Search system are tempted to rearrange the order of items in order to influence which gets paired with which. The purpose of this paper is to explain in more detail how the system works, and why it does not matter what the actual pairings are.

It is often asked how the computer 'decides' which item to pair with which - is it random? Well, it is and it isn't - we describe the pairings as *pseudorandom*. There is a numerical structure underlying the questionnaire - there has to be for the system to work at all. Which item is paired with which follows a set of numerical codes. For example, in a 38-item questionnaire, item 1 is always paired with items 20, 29 and 38. Item 5 happens to be paired with items 23, 24 and 33, and so on. Item number relates to the position of the item in the original list which was first typed into the machine, and since the operator is not aware of the structure of the code system, there can be no intentional relation between the meanings of the items. So in this sense, the pairings are random.

People are sometimes concerned that certain pairings, for example of a highly favoured item with three which are strongly disliked, will introduce bias. A similar concern relates to what happens if the respondent marks the middle of the line on each of the three occasions the item appears. Surely we have no information at all about that item? To see why these are not problematical needs more insight into how the paired comparison system actually works. The important thing to understand is that none of the pairings exists in isolation - each is part of an extensive network which involves *all* of the other items:

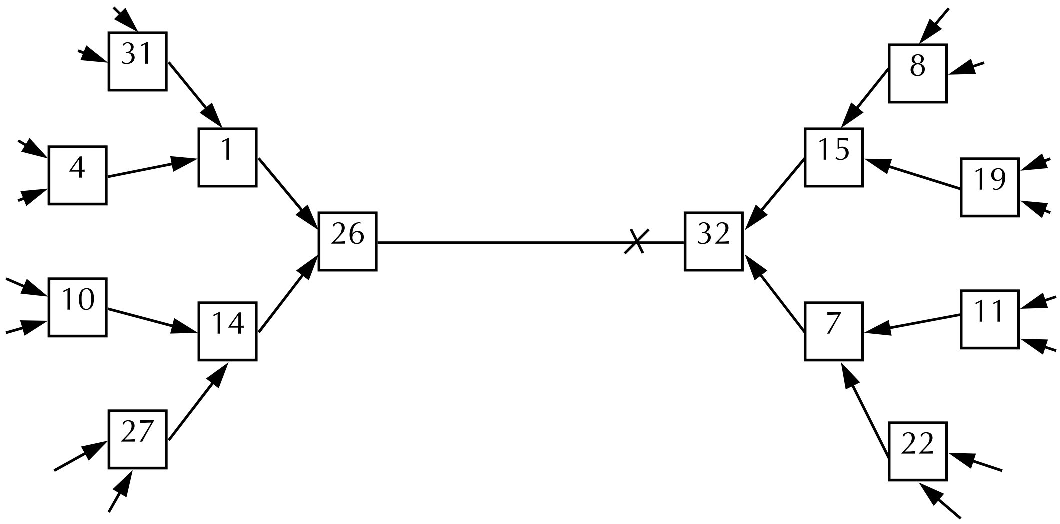

The mark on the line between items 26 and 32 favours 32; the strength of this weighting in favour of item 32, however, is influenced by the actual importance of item 26. This in turn derives from not only from the weights it receives relative to items 1 and 14, but the actual importance of these items, which in turn.... So you see that the process of calculating a rank order from paired comparisons is quite complex, and the result is that the effects of any idiosyncratic pairings are smoothed out.

Priority Research

In theory, it doesn't matter at all which item gets paired with which. If you create a questionnaire where the items are random numbers, and the mark on the line accurately reflects the relative sizes of the numbers, the system will always sort the numbers into order of size, quite regardless of what happens to be paired with what. It makes no difference if a large number happens to get paired with three small ones, it still ends up in its proper position in the list. The system is so effective, in fact, that with precise information like this it will sort random numbers correctly using only two pairings, instead of the three which the system actually uses.

In practice, matters are somewhat different. A measure of how much I like something can never be as precise as a measure of how heavy I am, or how hot the day is. People are not consistent, they do not rate their preferences with mathematical precision, and they may be unduly influenced by items which they see as particularly important. It is thus possible in the real world for the accident of which item happens to be paired with which to affect the priority list which emerges. However, this tendency is completely swamped by a much larger source of error, and one which is unavoidable.

If somebody fills in a Priority Search questionnaire twice, the two preference orders produced will be different - similar, but not identical. The more items there are in the questionnaire, the more scope there will be for any individual item to shift position somewhat. Whilst you will never find that the top item shifts to the bottom, or vice versa, movements of two or three positions either way are normal in a questionnaire with, say, 30 items. One or two shifts will be larger than this. It turns out that the variability, or, if you like, uncertainty, in the rank order of items for each individual caused by this effect is *three times* the size of the variability introduced by the pairing of particular items.

This fluctuation in response to the same questionnaire by the same individual is not a defect in the system, it is inherent in the way people ascribe value (i.e. by placing a mark on a line) to their subjective feelings. Fortunately, this fluctuation is random, and when the preferences of a group are displayed (which is how the system is normally used) the fluctuations tend to cancel each other out. Thus, whilst an *individual* will give somewhat different rank orders on different occasions, a *group* will be much more consistent. All of the statistical tests used to calculate significances take this variability into account.

People are complicated things, and measuring preferences can never be a 'hard' science. Nevertheless, the Priority Search system for detecting people's priorities is extremely effective. Tests have shown that it is at least as reliable as the best questionnaires used for psychometric testing, and better than most. Its use for over thirteen years in many hundreds of projects, in widely differing fields, has made it the market leader in this kind of research.