Choosing the right order for sequential treatments on cost-effectiveness grounds: the case of rheumatoid arthritis

Pelham Barton, Paresh Jobanputra, Stirling Bryan, Amanda Burls
University of Birmingham

Author for Correspondence & Contact Details
Pelham Barton
Health Economics Facility
University of Birmingham
Park House
40 Edgbaston Park Road
Birmingham
B15 2RT
P.M.Barton@bham.ac.uk

Issue Being Addressed
Rheumatoid arthritis (RA) is a chronic condition for which a large number of disease-modifying anti-rheumatic drugs (DMARDs) are available. Typically DMARDs will be stopped after a time for reasons either of loss of efficacy or toxicity. Thus the appropriate long-term strategy for treating a patient with RA requires the use of a sequence of DMARDs. (DMARDs may also be combined in various ways.) Any possible sequence of DMARDs is in principle a candidate for cost-effectiveness analysis. With 11 commonly used DMARDs to consider, there are a total of nearly 40 million possible sequences of DMARDs, and over 100 million sequences if subsets are to be considered as well. It is clearly not possible to test all these sequences.

Methods
Sequences of DMARDs are compared using the Birmingham Rheumatoid Arthritis Model (BRAM), which was developed as part of the NICE technology appraisals programme. Features of the BRAM are that it allows realistic distributions to be used for the time spent on any DMARD, and that it can allow the effects of a DMARD to depend on a patient's previous history. The model is an individual sampling model, which works by generating a large number of virtual patient histories from which population mean costs and QALYs are estimated. Local search techniques from operational research are used to find the optimum sequence in a computationally feasible way. Uncertainty in data can be accounted for by deterministic or probabilistic sensitivity analysis: both approaches are discussed in this paper.

Results
Optimal sequences of DMARDs are shown under a variety of modelling assumptions.

Key Issues for Discussion by HESG Audience
The appropriate use and presentation of local search techniques in the context of cost-effectiveness analysis.

Scope of this Paper

This paper represents work in progress, and should not be cited or quoted without permission from the authors. The focus of this paper is methodological: we are concerned with the use of local search techniques to find an optimum strategy from an infeasibly large number of options. The underlying model used for this analysis is the Birmingham Rheumatoid Arthritis Model (BRAM), which has been used for a previous NICE appraisal (Clark *et al*, 2004). The BRAM has been criticised in the past, (for example, Bansback *et al*, 2005), principally on the basis that our critics, who have contributed to models developed for the pharmaceutical industry, would prefer us to use different numerical inputs into the model. We recognise that there is scope for improving the BRAM model structure and are doing this as part of an ongoing NICE appraisal (NICE, 2005). Such development is subject to the usual reporting restrictions for an appraisal in progress, and therefore the work reported here is based on the data used in the currently published version of the model (Clark *et al*, 2004).

Background

Rheumatoid arthritis (RA) is a systemic inflammatory disorder that mainly affects synovial joints. RA is characterised by pain, swelling and stiffness of synovial joints. Conventional drug therapy for RA relies on varying combinations of the following four classes of drugs:

- non-steroidal anti-inflammatory drugs (NSAIDs);
- analgesics;
- corticosteroids such as prednisolone and methylprednisolone;
- disease-modifying antirheumatic drugs (DMARDs).

Conventional DMARDs are slow-acting drugs, taking weeks or months to provide symptomatic relief and reduce the risk of progressive joint damage. A patient taking a particular DMARD may have to stop because of either toxicity, loss of efficacy, patient or physician preference or development of disease complications or other co-morbidity. DMARDs may be used sequentially or in combination. For the purpose of this paper, only sequential use will be considered. Table 1 shows the DMARDs used in the BRAM and some of their properties, summarised from Clark *et al* (2004). Costs shown here exclude additional costs when starting a new DMARD.

Table 1 DMARDs used in this work

DMARD	Symbol	Annual cost (2002 £)	QoL gain while effective	Mean time to failure (yrs)
Anakinra	Ana	8080.32	0.082	1.77
Azathioprine	AZA	1286.78	0.082	1.95
Ciclosporin A	CyA	2963.63	0.123	8.71
Etanercept	Etan	9513.64	0.164	15.03
Gold (injectable)	GST	1450.08	0.123	3.85
Hydroxychloroquine	HCQ	426.74	0.082	3.62
Infliximab	Infl	9867.24	0.164	7.26
Leflunomide	LEF	1124.60	0.123	2.28
Methotrexate	MTX	1156.37	0.164	5.39
Penicillamine	DPen	1334.81	0.082	2.69
Sulfasalazine	SSZ	510.10	0.123	3.45

Any possible sequence of DMARDs is in principle a candidate for cost-effectiveness analysis. With 11 commonly used DMARDs to consider, there are a total of nearly 40 million possible sequences of DMARDs, and over 100 million sequences if subsets are to be considered as well. It is clearly not possible to test all these sequences.

Basic Model Structure

Full details of the BRAM are in the two HTA reports Barton *et al* (2004b) and Clark *et al* (2004). The BRAM is an individual sampling model (Barton *et al*, 2004a). When the model is run to compare two strategies, a large number of individuals are run through each strategy in turn. Lifetime (discounted) costs and QALYs are accumulated for each individual, and, in this paper, have been converted into net monetary benefit (NMB), using a threshold ICER of £30,000/QALY in the base case.

An important feature of the model is that it acknowledges the correct decision point as the point of divergence between strategies in terms of the drug used next in the sequence. Consider, for example, the first comparison in Clark *et al* (2004). The drug sequences compared are as follows:

Baseline: (SSZ – MTX – LEF – Etan – Infl –) GST – AZA – CyA
Comparator: (SSZ – MTX – LEF – Etan – Infl –) Ana – GST – AZA – CyA
(In fact, each sequence continues with combination therapy of MTX+CyA, but that is unimportant to this discussion.)

Effectively the decision to be evaluated here relates to patients who have failed the first five drugs common to the two sequences, shown in brackets above. For such patients, the question is whether they should immediately receive injectable gold

(GST) or first be given anakinra (Ana), with the intention of following the rest of the baseline sequence after failure on anakinra.

One way of modelling this would be to define a starting population of people who have failed the first five (common) drugs, and use the model to compare the three or four remaining drugs in the rest of the sequences. If this were done, discounting would automatically and appropriately be to the divergence point between the strategies. If a model were to be built solely for the purpose of this single comparison, this would undoubtedly be the preferred approach.

The disadvantage of that method comes when a range of different sequences are to be compared. For example, Clark *et al* also considered the above comparison but with sequences that did not feature etanercept. To do this, using the method suggested in the last paragraph, would mean constructing a different starting population to which to apply the strategies.

The method used in the BRAM is always to start all patients from the beginning of the sequence of DMARDs. Those patients who remain for sufficiently long on early DMARDs, and who, thus, do not reach the divergence point, are discarded, effectively creating an appropriate starting population within the model. For those who reach the divergence point, only costs and QALYs after the divergence point are counted, and these are discounted back to the divergence point and not back to the beginning of the sequence.

Thus, within the model, a strategy is defined by a sequence of drugs and a divergence point. The only meaningful comparisons are between sequences which are the same up to the divergence point, but then vary. When considering adding a new drug to the end of a sequence, or removing the last drug from the existing sequence, the divergence point on the shorter sequence is simply at the end of the shorter sequence of DMARDs, when the patient moves to palliation.

Determining the Better Sequence

An efficiency improvement has been made to the model reported in Clark *et al* (2004) and so the model runs reported in this paper used what we will describe as the "HESG version". The "HESG version" of the model runs in the following way to compare two strategies. Virtual patients are run through the model from the start of each strategy and first DMARD common to the two strategies. Any patient who does not reach the divergence point is discarded. When a patient reaches the divergence point,

the patient's characteristics are stored and the patient is then run through the remainder of each strategy in turn. In the case where the divergence point is at the start of the sequence, no patients are discarded and each newly generated patient is run through each of the strategies. For each individual patient there is then a difference in cost c_{diff} and a difference in QALYs e_{diff}. This is converted to net monetary benefit (NMB) using the formula

$$NMB = e_{diff}\lambda - c_{diff},$$

where λ is the threshold ICER (assumed to be £30,000 per QALY in the base case for the purposes of this paper). The mean NMB across a large number of patients is an approximation to the (modelled) population mean. To ensure that the sampling mean is a good approximation to the population mean, it is important to know if a sufficient number of replications of the model have been run. The variation among the results for the modelled individuals gives us a measure of the uncertainty of the mean due to the fact that the model is run only for a finite number of individuals. Following Briggs (2000), we refer to this as the "quasi standard error" (QSE). It is important to note that this does not reflect any sort of parameter or structural uncertainty, merely the sampling within the model for a fixed set of parameters.

For a single comparison, one could form an approximately 95% quasi confidence interval for the difference in NMB between two strategies by $\bar{x}_{diff} \pm 2s_{diff}$, where $\bar{x}_{diff}$ and s_{diff} are the sampled mean and its standard error, respectively. However, by definition, such intervals do not include the true mean approximately 5% of the time. Given the large number of comparisons to be made in applying the optimisation algorithm, the risk of incorrect deductions is not negligible. For the work reported below, we used wider quasi confidence intervals of $\bar{x}_{diff} \pm 4s_{diff}$. Only when both ends of this confidence interval had the same numerical sign did we regard one sequence as securely better than the other.

Basic Principle of the Descent Algorithm

The idea is to start with some possible sequence of DMARDs, then compare it with a similar ("neighbouring") sequence. Under suitable conditions, repeatedly replacing sequences by better sequences leads eventually to the optimal sequence. To apply the algorithm, it is necessary to define the concept of a neighbourhood for a given sequence. The essential property is that it must be possible to get from any sequence to any other sequence. This property may be called "connectedness". A desirable property is that neighbourhoods should be as small as possible, consistent with the aim of producing an optimum sequence. For complete sequences of all DMARDs, connectedness is satisfied by allowing exchange of any pair of DMARDs that are

adjacent in the sequence. This is generally known as “adjacent pairwise interchange (API)”. To allow for subsequences, we can also allow dropping of the last DMARD in the sequence, and adding to the end of the sequence any DMARD not currently included. The question of whether this neighbourhood is sufficient to produce an optimum sequence will be discussed later in this paper.

The Modified API Neighbourhood

Suppose that at some stage in the process, the current sequence of drugs is A – B – C – D – E – F – G, drugs H, I, J, and K having been excluded (H most recently). Then the following sequences are defined as neighbouring:

Comparison	Current Sequence	Neighbouring Sequence
Swap A with B	A – B – C – D – E – F – G	B – A – C – D – E – F – G
Swap B with C	(A –) B – C – D – E – F – G	(A –) C – B – D – E – F – G
Swap C with D	(A – B –) C – D – E – F – G	(A – B –) D – C – E – F – G
Swap D with E	(A – B – C –) D – E – F – G	(A – B – C –) E – D – F – G
Swap E with F	(A – B – C – D –) E – F – G	(A – B – C – D –) F – E – G
Swap F with G	(A – B – C – D – E –) F – G	(A – B – C – D – E –) G – F
Exclude G	(A – B – C – D – E – F –) G	(A – B – C – D – E – F)
Add H	(A – B – C – D – E – F – G)	(A – B – C – D – E – F – G –) H
Add I	(A – B – C – D – E – F – G)	(A – B – C – D – E – F – G –) I
Add J	(A – B – C – D – E – F – G)	(A – B – C – D – E – F – G –) J
Add K	(A – B – C – D – E – F – G)	(A – B – C – D – E – F – G –) K

Figure 1 shows the complete set of possibilities for sequences of three or fewer DMARDs. Neighbouring sequences are joined by lines. The significance of the heavier lines is discussed below.

Figure 1 Sequences of three or fewer DMARDs

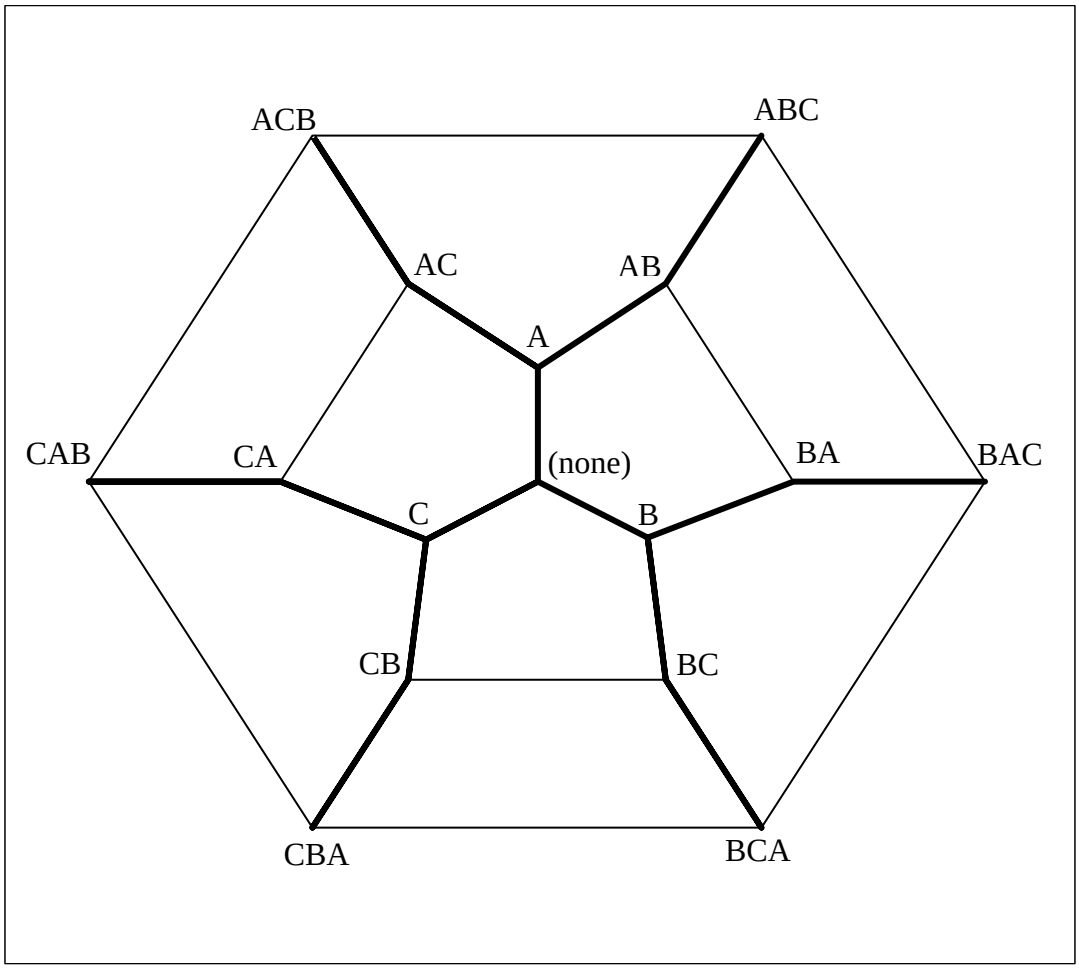

Sufficiency of the Modified API Neighbourhood

A potential problem with local search algorithms is the risk of the process stopping at a "local optimum", that is to say a sequence that is better than any of its neighbours, but is not the true optimum sequence. Informal consideration of the way the BRAM works suggests that this is highly unlikely to be a problem using the modified API neighbourhood as defined above. Consider the case of only three drugs illustrated in Figure 1 and suppose, for example, that the true optimum sequence is ABC. Then it can reasonably be expected that both CAB and BCA would be better than CBA, that ACB would be better than CAB, and that BAC would be better than BCA. Accordingly, no sequence of all three drugs can be a local optimum. Similar considerations apply to subsequences, and to cases where the true optimum is a subsequence of the three drugs.

The heavier lines in Figure 1 relate to neighbouring sequences defined by exclusion of the last DMARD (or adding a new DMARD to a subsequence). It can be seen on the diagram that this more restricted definition of neighbourhood still satisfies the connectedness property. However, supposing again that the true optimum strategy is to use all three drugs and that the best order is ABC, it is likely that any sequence of all three drugs is better than the strategy of using only the first two drugs in the sequence. In other words, each sequence of all three drugs (other than the true optimum) is likely to be a local optimum. Thus this more restricted definition of neighbourhood is not suitable for our purpose.

Further consideration will be given to the "local optimum" problem later in this paper. Having defined the basis for comparison between strategies and the choice of neighbourhood, we are now in a position to give a formal statement of the algorithm used to find an optimal sequence.

Optimisation Using the Descent Algorithm

The basic algorithm used to determine the optimal sequence of DMARDs is as follows:

Step 1 Select a starting sequence at random to be the current sequence.

Step 2 Set the number of patients to be tried for each sequence to 10,000.

Step 3 Try each neighbour of the current sequence in turn. If, in any of them, the comparator is securely better than the current sequence, then replace the current sequence by that comparator and return to step 2. If the current sequence is securely better than the comparator, then note that fact. If the current sequence is securely better than all comparators, then go to step 6. If any comparisons are not securely determined, and no comparator is securely better than the current sequence, then go to step 4.

Step 4 Increase the number of patients for each sequence to the next number from the list 20,000, 40,000, 100,000, 200,000, 400,000, 1,000,000, etc.

Step 5 Repeat step 3, omitting any comparisons where the current sequence has already been shown to be securely better than the comparator. Go to step 2, 4, or 6, as appropriate.

Step 6 Finish. The current sequence is a local or global optimum sequence.

Results

The process described above has been undertaken for the RA DMARD sequences, using inputs including reported in Clark *et al* (2004), some of which are shown in Table 1. The initial strategy selected at random was as follows:

Infl – AZA – DPen – SSZ – Ana – LEF – CyA – MTX – GST – Etan – HCQ

With 10,000 patients in each arm, the first secure improvement was to swap Etan with HCQ. The results were as shown below:

Strategy	Mean (QSE) cost (£)	Mean (QSE) QALY
Baseline	46824 (367)	4.4147 (0.0397)
Swap Etan with HCQ	33628 (326)	4.3016 (0.0395)
	Mean NB (£)	Q.S.E. (£)
Difference	9802	728

Accept change to

Infl – AZA – DPen – SSZ – Ana – LEF – CyA – MTX – GST – HCQ – Etan

This time, the only secure change was to drop Etan.

Strategy	Mean (QSE) cost (£)	Mean (QSE) QALY
Baseline	45698 (357)	3.9435 (0.0370)
Drop Etan	2627 (15)	2.9405 (0.0321)
	Mean NB (£)	Q.S.E. (£)
Difference	12980	672

Accept change to

Infl – AZA – DPen – SSZ – Ana – LEF – CyA – MTX – GST – HCQ

This time, the first secure change was to swap CyA with MTX.

Strategy	Mean (QSE) cost (£)	Mean (QSE) QALY
Baseline	18421 (96)	6.2486 (0.0463)
Drop Etan	16090 (90)	6.3735 (0.0462)
	Mean NB (£)	Q.S.E. (£)
Difference	6079	1036

Accept change to

Infl – AZA – DPen – SSZ – Ana – LEF – MTX – CyA – GST – HCQ

Continuing this process, changes to the following sequences in turn were accepted:

AZA – Infl – DPen – SSZ – Ana – LEF – MTX – CyA – GST – HCQ
AZA – DPen – Infl – SSZ – Ana – LEF – MTX – CyA – GST – HCQ
AZA – DPen – SSZ – Infl – Ana – LEF – MTX – CyA – GST – HCQ

At this stage, no secure improvements were found at 10,000 patients in each arm, but the options "swap SSZ with Infl", "drop HCQ", and "add Etan" were securely rejected. Re-running the remaining eight options with 20,000 patients in each arm still led to no secure improvements, but the option "swap MTX with CyA" was securely rejected. The remaining sveen options still led to no secure improvement at 40,000 patients per arm, but "swap Ana with LEF" was accepted at 100,000 patients in each arm. Thus the current sequence was changed to

AZA – DPen – SSZ – Infl – LEF – Ana – MTX – CyA – GST – HCQ

and then successively to the following sequences:

AZA – DPen – SSZ – LEF – Infl – Ana – MTX – CyA – GST – HCQ
AZA – DPen – SSZ – LEF – Infl – MTX – Ana – CyA – GST – HCQ
AZA – DPen – SSZ – LEF – MTX – Infl – Ana – CyA – GST – HCQ
AZA – DPen – SSZ – LEF – MTX – Infl – CyA – Ana – GST – HCQ
AZA – DPen – SSZ – LEF – MTX – CyA – Infl – Ana – GST – HCQ
AZA – DPen – SSZ – LEF – MTX – CyA – Infl – GST – Ana – HCQ
AZA – DPen – SSZ – LEF – MTX – CyA – GST – Infl – Ana – HCQ
AZA – DPen – SSZ – LEF – MTX – CyA – GST – Infl – HCQ – Ana
AZA – DPen – SSZ – LEF – MTX – CyA – GST – HCQ – Infl – Ana
AZA – DPen – SSZ – LEF – MTX – CyA – GST – HCQ – Infl
AZA – DPen – SSZ – LEF – MTX – CyA – GST – HCQ
AZA – DPen – SSZ – LEF – MTX – CyA – HCQ – GST
AZA – DPen – SSZ – LEF – MTX – HCQ – CyA – GST
AZA – DPen – SSZ – LEF – MTX – HCQ – GST – CyA
AZA – SSZ – DPen – LEF – MTX – HCQ – GST – CyA
AZA – SSZ – DPen – MTX – LEF – HCQ – GST – CyA
AZA – SSZ – MTX – DPen – LEF – HCQ – GST – CyA
SSZ – AZA – MTX – DPen – LEF – HCQ – GST – CyA
SSZ – MTX – AZA – DPen – LEF – HCQ – GST – CyA
SSZ – MTX – AZA – DPen – HCQ – LEF – GST – CyA
SSZ – MTX – AZA – HCQ – DPen – LEF – GST – CyA
SSZ – MTX – HCQ – AZA – DPen – LEF – GST – CyA
SSZ – MTX – HCQ – AZA – LEF – DPen – GST – CyA
SSZ – MTX – HCQ – LEF – AZA – DPen – GST – CyA
SSZ – MTX – HCQ – LEF – DPen – AZA – GST – CyA
SSZ – MTX – HCQ – LEF – DPen – GST – AZA – CyA
SSZ – MTX – HCQ – LEF – GST – DPen – AZA – CyA
SSZ – MTX – HCQ – LEF – GST – DPen – CyA – AZA

When the model was run from the final sequence above, all changes were rejected and the final sequence is thus a local or global optimum.

The model was re-run starting with the empty sequence of DMARDs. This time the changes made were as shown below:

(none)
SSZ
SSZ – CyA
SSZ – CyA – LEF
SSZ – CyA – LEF – MTX
SSZ – CyA – LEF – MTX – GST
SSZ – CyA – LEF – MTX – GST – HCQ
SSZ – CyA – LEF – MTX – HCQ – GST
SSZ – CyA – LEF – MTX – HCQ – GST – DPen
SSZ – CyA – LEF – MTX – HCQ – GST – DPen – AZA
SSZ – LEF – CyA – MTX – HCQ – GST – DPen – AZA
SSZ – LEF – MTX – CyA – HCQ – GST – DPen – AZA
SSZ – LEF – MTX – HCQ – CyA – GST – DPen – AZA
SSZ – LEF – MTX – HCQ – GST – CyA – DPen – AZA
SSZ – MTX – LEF – HCQ – GST – CyA – DPen – AZA
SSZ – MTX – HCQ – LEF – GST – CyA – DPen – AZA
SSZ – MTX – HCQ – LEF – GST – DPen – CyA – AZA

Final sequence accepted as a local or global optimum.

It can be seen that the process has converged to the same final sequence from two different starting points. It can also be noted that once a pair of DMARDs had changed places, they never changed back again. It is thus reasonable to conclude that the final sequence reached is a true optimum solution of the model. Indeed, it is relevant that the derived sequence is clinically credible (Jobanputra *et al*, 2004). Where there is reasonable doubt as to whether a true optimum has been reached, methods can be used such as starting again from a different random sequence, or testing the final sequence against an enlarged neighbourhood.

Of course, the results above relate to a specific parameter set, and it is appropriate to consider the effects of uncertainty in those parameters. Before doing so, we shall consider variations on the descent algorithm itself. A good survey of local search methods was written by Colorni *et al* (1996).

Variations on the Descent Algorithm

One possible variation would be to consider each neighbour of the current sequence in turn, and make a secure decision on that possibility before going on to the next. It would also be possible to consider the neighbouring sequences in a random order. The drawback with this approach is that on occasions it requires a large number of replications to decide not to make a particular change: this computing time would be effectively wasted, compared to the method used, which aims to find a positive movement as quickly as possible.

Another possibility is to find which neighbour of the current sequence gives the greatest improvement in NMB, and change to that one. However, consider moving towards the optimum sequence from some substantially different order of the same DMARDs. Experience with the model, confirmed by the results shown above, suggests that the only possible local improvements to any sequence are to swap a pair of (adjacent) DMARDs whose order in the current sequence is the reverse of their order in the optimal sequence. It is not hard to see that all possible routes from the current sequence to the optimal sequence involve the same number of swaps. For example, consider four DMARDs whose optimal sequence is ABCD and start from DCBA (as far away as possible). Then the following routes may be used to reach the optimum sequence:

DCBA – CDBA – CBDA – BCDA – BCAD – BACD – ABCD
DCBA – CDBA – CBDA – CBAD – BCAD – BACD – ABCD
DCBA – CDBA – CBDA – CBAD – CABD – ACBD – ABCD
DCBA – CDBA – CDAB – CADB – ACDB – ACBD – ABCD
DCBA – CDBA – CDAB – CADB – CABD – ACBD – ABCD
DCBA – DBCA – BDCA – BCDA – BCAD – BACD – ABCD
DCBA – DBCA – BDCA – BDAC – BADC – ABDC – ABCD
DCBA – DBCA – BDCA – BDAC – BADC – BACD – ABCD
DCBA – DBCA – DBAC – BDAC – BADC – ABDC – ABCD
DCBA – DBCA – DBAC – BDAC – BADC – BACD – ABCD
DCBA – DBCA – DBAC – DABC – ADBC – ABDC – ABCD
DCBA – DCAB – CDAB – CADB – ACDB – ACBD – ABCD
DCBA – DCAB – CDAB – CADB – CABD – ACBD – ABCD
DCBA – DCAB – DACB – ADCB – ACDB – ACBD – ABCD
DCBA – DCAB – DACB – ADCB – ADBC – ABDC – ABCD
DCBA – DCAB – DACB – DABC – ADBC – ABDC – ABCD

As can be seen, each of the sixteen possible routes consists of six swaps. This confirms that there is no obvious advantage in trying to make the largest improvement

first. In any case, the system of selecting a change that is securely accepted on the smallest number of replications of the model tends to favour making larger improvements first.

Tabu Search

(*The spelling "tabu" is used as standard in the literature in preference to the more normal British spelling "taboo". In either case, the British pronunciation puts the stress on the second syllable.*)

The searching process used in the descent algorithm could be made more efficient if moves to neighbouring sequences were not considered if there was good reason to suppose in advance that they would be rejected. An obvious example of this is the immediate return to the previous "current sequence". As was observed earlier, once any pair of DMARDs had been swapped, they were never swapped back. This suggests that it is unnecessary to consider such swaps. Similarly, once a DMARD had been dropped from the sequence, it was never reinstated.

The argument can be taken further. If a change is securely rejected, then it can be assumed that a similar change will not be beneficial in future iterations of the descent algorithm.

The process of tabu search is similar to the descent algorithm, except that there is a "tabu list" of changes which will not be considered in future iterations of the model. In the case of the BRAM, these could include:

- reversal of swaps actually made;
- reinstatement of DMARDs dropped (if starting with a list of all DMARDs);
- dropping of DMARDs added (if starting "empty");

and possibly

- swaps rejected in previous iterations;
- adding DMARDs whose addition was rejected in previous iterations;
- dropping DMARDs if this was rejected in previous iterations.

If this had been applied to the BRAM as shown above, the results reported would be exactly the same. This suggests that in this case it would be reasonable to use a tabu list, and to make it permanent. In other situations, it may be preferable to make the entries on a tabu list temporary. This can be done in one of (at least) two ways. One way is to put an explicit "expiry date" on any entry, for example, by stating that a tabu lasts five iterations of the algorithm. An alternative is to limit the length of the tabu list: once the list has reached its full length, new entries displace the oldest existing entries. In any case, the final sequence reached should be tested against (at least) the entire neighbourhood used for the search.

Sensitivity Analysis

Of course, it is necessary to take account of uncertainty in the model parameters. This can be done through deterministic or probabilistic sensitivity analysis.

Deterministic sensitivity analysis involves re-running the descent algorithm with a different parameter set. The first illustration involves re-running the model with a threshold ICER of £100,000/QALY (instead of £30,000/QALY). Starting with an empty sequence, the model made the following changes:

(none)
Infl
Infl – SSZ
Infl – SSZ – DPen
Infl – SSZ – DPen – CyA
Infl – SSZ – DPen – CyA – LEF
Infl – SSZ – DPen – CyA – LEF – MTX
Infl – SSZ – DPen – CyA – LEF – MTX – GST
Infl – SSZ – DPen – CyA – LEF – MTX – GST – AZA
Infl – SSZ – DPen – CyA – LEF – MTX – GST – AZA – Etan
Infl – SSZ – DPen – CyA – LEF – MTX – GST – AZA – Etan – HCQ
SSZ – Infl – DPen – CyA – LEF – MTX – GST – AZA – Etan – HCQ
SSZ – DPen – Infl – CyA – LEF – MTX – GST – AZA – Etan – HCQ
SSZ – DPen – CyA – Infl – LEF – MTX – GST – AZA – Etan – HCQ
SSZ – DPen – CyA – LEF – Infl – MTX – GST – AZA – Etan – HCQ
SSZ – DPen – CyA – LEF – MTX – Infl – GST – AZA – Etan – HCQ
SSZ – DPen – CyA – LEF – MTX – GST – Infl – AZA – Etan – HCQ
SSZ – DPen – CyA – MTX – LEF – GST – Infl – AZA – Etan – HCQ
SSZ – DPen – MTX – CyA – LEF – GST – Infl – AZA – Etan – HCQ
SSZ – MTX – DPen – CyA – LEF – GST – Infl – AZA – Etan – HCQ
SSZ – MTX – CyA – DPen – LEF – GST – Infl – AZA – Etan – HCQ
SSZ – MTX – CyA – DPen – LEF – GST – Infl – AZA – HCQ – Etan

SSZ – MTX – CyA – DPen – LEF – GST – Infl – HCQ – AZA – Etan
SSZ – MTX – CyA – DPen – LEF – GST – HCQ – Infl – AZA – Etan
SSZ – MTX – CyA – LEF – DPen – GST – HCQ – Infl – AZA – Etan
SSZ – MTX – CyA – LEF – DPen – GST – HCQ – Infl – Etan – AZA
SSZ – MTX – CyA – LEF – GST – DPen – HCQ – Infl – Etan – AZA
SSZ – MTX – CyA – LEF – GST – HCQ – DPen – Infl – Etan – AZA
SSZ – MTX – CyA – LEF – GST – HCQ – DPen – Etan – Infl – AZA
SSZ – MTX – CyA – LEF – GST – HCQ – Etan – DPen – Infl – AZA
MTX – SSZ – CyA – LEF – GST – HCQ – Etan – DPen – Infl – AZA
MTX – SSZ – CyA – LEF – GST – HCQ – Etan – DPen – AZA – Infl
MTX – SSZ – CyA – LEF – GST – HCQ – Etan – DPen – AZA – Infl
MTX – SSZ – LEF – CyA – GST – HCQ – Etan – DPen – AZA – Infl
MTX – SSZ – LEF – GST – CyA – HCQ – Etan – DPen – AZA – Infl
(final sequence not confirmed as optimal – see discussion later)

To illustrate the effect of changes in model parameters, it has been claimed that the biologics (anakinra, etanercept, infliximab) slow the general decline in patients' quality of life. To see the importance of this claim, the model was re-run with virtually no decline in patients' quality of life while on these three drugs. Again starting from the empty sequence, the model gives the following changes with a threshold ICER of £30,000/QALY:

(none)
Infl
Infl – SSZ
Infl – SSZ – DPen
Infl – SSZ – DPen – CyA
Infl – SSZ – DPen – CyA – MTX
Infl – SSZ – DPen – CyA – MTX – GST
Infl – SSZ – DPen – CyA – MTX – GST – LEF
Infl – SSZ – DPen – CyA – MTX – GST – LEF – Etan
Infl – SSZ – DPen – CyA – MTX – GST – LEF – Etan – HCQ
Infl – SSZ – DPen – MTX – CyA – GST – LEF – Etan – HCQ
Infl – SSZ – DPen – MTX – CyA – GST – LEF – HCQ – Etan
SSZ – Infl – DPen – MTX – CyA – GST – LEF – HCQ – Etan
SSZ – Infl – DPen – MTX – CyA – GST – LEF – HCQ – Etan – AZA
SSZ – DPen – Infl – MTX – CyA – GST – LEF – HCQ – Etan – AZA
SSZ – DPen – MTX – Infl – CyA – GST – LEF – HCQ – Etan – AZA
SSZ – DPen – MTX – CyA – Infl – GST – LEF – HCQ – Etan – AZA
SSZ – DPen – MTX – CyA – GST – Infl – LEF – HCQ – Etan – AZA
SSZ – DPen – MTX – CyA – GST – LEF – Infl – HCQ – Etan – AZA
SSZ – DPen – MTX – CyA – GST – LEF – HCQ – Infl – Etan – AZA
SSZ – DPen – MTX – CyA – GST – LEF – HCQ – Etan – Infl – AZA
SSZ – DPen – MTX – CyA – GST – LEF – HCQ – Etan – AZA – Infl
SSZ – DPen – MTX – CyA – GST – LEF – HCQ – AZA – Etan – Infl
SSZ – MTX – DPen – CyA – GST – LEF – HCQ – AZA – Etan – Infl
SSZ – MTX – DPen – GST – CyA – LEF – HCQ – AZA – Etan – Infl

SSZ – MTX – DPen – GST – LEF – CyA – HCQ – AZA – Etan – Infl
SSZ – MTX – DPen – GST – LEF – HCQ – CyA – AZA – Etan – Infl
SSZ – MTX – DPen – LEF – GST – HCQ – CyA – AZA – Etan – Infl
SSZ – MTX – DPen – LEF – HCQ – GST – CyA – AZA – Etan – Infl
SSZ – MTX – DPen – HCQ – LEF – GST – CyA – AZA – Etan – Infl
SSZ – MTX – HCQ – DPen – LEF – GST – CyA – AZA – Etan – Infl
SSZ – MTX – HCQ – LEF – DPen – GST – CyA – AZA – Etan – Infl
SSZ – MTX – HCQ – LEF – GST – DPen – CyA – AZA – Etan – Infl
(final sequence not confirmed as optimal – see discussion later)

Using a threshold of £100,000/QALY with no decline in quality of life on biologics gave the following:

(none)
Infl
Infl – SSZ
Infl – SSZ – AZA
Infl – SSZ – AZA – Ana
Infl – SSZ – AZA – Ana – CyA
Infl – SSZ – AZA – Ana – CyA – LEF
Infl – SSZ – AZA – Ana – CyA – LEF – MTX
Infl – SSZ – AZA – Ana – CyA – LEF – MTX – GST
Infl – SSZ – AZA – Ana – CyA – LEF – MTX – GST – DPen
Infl – SSZ – AZA – Ana – CyA – LEF – MTX – GST – DPen – Etan
Infl – SSZ – AZA – Ana – CyA – LEF – MTX – GST – DPen – Etan – HCQ
Infl – SSZ – AZA – Ana – CyA – LEF – MTX – GST – Etan – DPen – HCQ
Infl – SSZ – AZA – Ana – CyA – LEF – MTX – Etan – GST – DPen – HCQ
Infl – SSZ – AZA – Ana – CyA – LEF – Etan – MTX – GST – DPen – HCQ
Infl – SSZ – AZA – Ana – CyA – Etan – LEF – MTX – GST – DPen – HCQ
Infl – SSZ – AZA – Ana – Etan – CyA – LEF – MTX – GST – DPen – HCQ
Infl – SSZ – AZA – Etan – Ana – CyA – LEF – MTX – GST – DPen – HCQ
Infl – SSZ – Etan – AZA – Ana – CyA – LEF – MTX – GST – DPen – HCQ
Infl – Etan – SSZ – AZA – Ana – CyA – LEF – MTX – GST – DPen – HCQ
Infl – Etan – SSZ – AZA – Ana – CyA – MTX – LEF – GST – DPen – HCQ
Infl – Etan – SSZ – AZA – Ana – MTX – CyA – LEF – GST – DPen – HCQ
Etan – Infl – SSZ – AZA – Ana – MTX – CyA – LEF – GST – DPen – HCQ
Etan – Infl – SSZ – Ana – AZA – MTX – CyA – LEF – GST – DPen – HCQ
Etan – Infl – SSZ – Ana – MTX – AZA – CyA – LEF – GST – DPen – HCQ
Etan – Infl – SSZ – Ana – MTX – CyA – AZA – LEF – GST – DPen – HCQ
Etan – Infl – Ana – SSZ – MTX – CyA – AZA – LEF – GST – DPen – HCQ
Etan – Infl – Ana – MTX – SSZ – CyA – AZA – LEF – GST – DPen – HCQ
Etan – Infl – Ana – MTX – SSZ – CyA – LEF – AZA – GST – DPen – HCQ
Etan – Infl – Ana – MTX – SSZ – CyA – LEF – GST – AZA – DPen – HCQ
Etan – Infl – Ana – MTX – SSZ – CyA – LEF – GST – AZA – HCQ – DPen
Etan – Infl – Ana – MTX – SSZ – CyA – LEF – GST – HCQ – AZA – DPen
Etan – Infl – Ana – MTX – SSZ – LEF – CyA – GST – HCQ – AZA – DPen
Etan – Infl – Ana – MTX – SSZ – LEF – CyA – GST – HCQ – DPen – AZA
(final sequence not confirmed as optimal – see discussion later)

It can be seen that quite different sequences are appearing in each case.

Probabilistic Sensitivity Analysis

Probabilistic sensitivity analysis (PSA) involves defining a joint probability distribution for the complete parameter set of a model. The model can then be run with parameter sets drawn from this distribution. Two main benefits of probabilistic sensitivity analysis have been claimed (see, for example, Claxton *et al*, 2005). The first relates to correct estimation of the expected costs and utilities of each option to be compared and the second to representation of the uncertainty.

In relation to expected costs and utilities, if there is a non-linear relationship between parameter values and model outputs (as is certainly the case with the BRAM), the expected value from PSA will differ from the results of a model applying mean values to the parameters separately. The methods described in this paper can easily be adapted to use PSA for this purpose by sampling from the parameter distribution each time a new patient is started.

As to representation of the uncertainty, advocates of probabilistic sensitivity analysis recommend the cost-effectiveness acceptability curve and cost-effectiveness frontier (Fenwick *et al*, 2001). However, the number of different strategies that can be handled by such curves is limited. It is simply not feasible to produce an accurate cost-effectiveness acceptability curve or frontier when there are 100 million options to be considered.

Computational Feasibility

When the descent algorithm is started from a substantially non-optimal sequence, the model finds beneficial changes quite quickly. As the algorithm continues, the number of virtual patients needed to produce secure improvements (or securely reject possible changes) increases until very large numbers are needed. A feature of individual sampling models is that, as well as giving estimates of the mean difference, they also, automatically, give an indication of the variability in outcomes among individual patients. (The same effect can be obtained from state-transition models by running them using first-order Monte Carlo simulation.) When the difference between two options is very close to the threshold, the amount of variability in population mean

that can be allowed without crossing that threshold is very small. Thus large numbers of replications of the model are needed. This is why the sensitivity analysis results have not (yet) been taken through to completion of the descent algorithm.

It may be thought that cohort models (such as Markov models) do not suffer from this problem. However, there is an issue relating to cycle length. State transition models make approximations by working in time cycles, under the assumption that modelling errors due to choice of cycle length are negligible. If the difference in modelled outcome between two options is quite small, this assumption may fail.

Interpretation of Model Results

Although in the base case the descent algorithm has been run through to completion, it is not claimed that the final sequence produced is necessarily optimal in practice. Rather, it is claimed that the sequence is a sensible basis from which to choose a sequence of treatments in discussion with patients.

Generalisability of the Methods Described in this Paper

In principle, the methods used in this paper can be applied to any model that considers management of a chronic condition by sequential treatment of single drugs. It has been assumed in the work above that there are no restrictions on the order in which drugs may be used. Restrictions could be considered, for example that a particular drug cannot be used before a specified other drug, or before a certain position in the sequence. Either of these could be easily accommodated.

The methods can be extended to combination therapy by treating combinations as other treatments. One complication that arose with the BRAM is that combination therapy would only be used after at least one of the drugs making up the combination had been used singly, and then would not be used at all if that drug had been quit on grounds of toxicity. However, with careful handling, such combinations could be modelled.

An important issue to consider is the adequacy of the neighbourhood structure. For example, there could be a drug (X, say) which must be used either early in the sequence or not at all. If the basis for this were reflected in the model parameters it

would not be possible to move from a long sequence without drug X to one in which X was used early. This is because X would have to be added at the end and then shifted into its correct place. However, the model parameters would be such that adding X to a long sequence would not represent an improvement. The problem could be overcome by allowing new drugs to be added at any point in the sequence, not just at the end, at the expense of a larger number of comparisons to be made at each stage of the process.

Conclusions

It has been shown how an optimisation algorithm from operational research can be applied to a complex model for the management of a chronic condition. It has also been shown that deterministic sensitivity analysis can be applied, and that probabilistic sensitivity analysis can be used to estimate the expected outcome allowing for uncertainty in the model parameters. The methods reported in this paper relate more to choice of appropriate treatment for a patient than appraisal of specific drugs, and thus fit more with the clinical guideline part of NICE's work programme than with technology appraisal. This study highlights the potential tension between the two main parts of NICE's work programme.

It is not claimed that these methods lead to a provably optimal strategy for management of a chronic condition. Differences in net monetary benefit between similar strategies are at times very small, and only become statistically significant with very large numbers of patients. (This is true even at the 95% level rather than the more stringent criterion of "secure improvement" used in this paper.)

The example in this paper has selected from approximately 100 million possible strategies for the treatment of RA patients. Even this is only a tiny fraction of the number of strategies that are possible in practice, allowing for combination therapy and other types of treatment. Even if the model ran effectively instantaneously, it would not be possible to produce sensible cost-effectiveness acceptability curves and frontiers. This is a case where it is unrealistic to expect to quantify all the uncertainty involved in decision making, and any claim to have done so is necessarily at best self-deception.

References

Bansback NJ, Young A, Brennan A (2005) The NICE reappraisal of biologics in 2005: what rheumatologists need to know (editorial). *Rheumatology* 44, 3-4.

Barton P, Bryan S, Robinson S (2004a) Modelling in the economic evaluation of health care: selecting the appropriate approach. *Journal of Health Services Research and Policy* 9(2), 110-118.

Barton P, Jobanputra P, Wilson J, Bryan S, Burls A (2004b) The use of modelling to evaluate new drugs for patients with a chronic condition: the case of antibodies against tumour necrosis factor in rheumatoid arthritis. *Health Technology Assessment* 8(11).

Briggs AH (2000) Handling uncertainty in cost-effectiveness models. *Pharmacoeconomics* 17, 479-500.

Clark W, Jobanputra P, Barton P, Burls A (2004) The clinical and cost-effectiveness of anakinra for the treatment of rheumatoid arthritis in adults: a systematic review and economic analysis. *Health Technology Assessment* 8(18).

Claxton K, Sculpher M, McCabe C, *et al* (2005) Probabilistic sensitivity analysis for NICE technology assessment: not an optional extra. *Health Economics* 14, 339-347.

Colorni A, Dorigo M, Maffioli F, Maniezzo V, Righini G, Trubian M (1996) Heuristics from nature for hard combinatorial optimization problems. *International Transactions in Operational Research* 3(1), 1-21.

Fenwick E, Claxton K, Sculpher M (2001) Representing uncertainty: the role of cost-effectiveness acceptability curves. *Health Economics* 10, 779-789.

Jobanputra P, Wilson J, Douglas K, Burls A (2004) A survey of British Rheumatologists DMARD preferences for rheumatoid arthritis. *Rheumatology* 43, 206-10.

NICE (2005) http://www.nice.org.uk/page.aspx?o=207026 (accessed 23 May 2005).