

THE ANALYTIC HIERARCHY PROCESS (AHP)

INTRODUCTION

The Analytic Hierarchy Process (AHP) is due to Saaty (1980) and is often referred to, eponymously, as the Saaty method. It is popular and widely used, especially in military analysis, though it is not, by any stretch of the imagination, restricted to military problems. In fact, in his book, (which is not for the mathematically faint of heart!) Saaty describes case applications ranging from the choice of a school for his son, through to the planning of transportation systems for the Sudan. There is much more to the AHP than we have space for but we will cover the most easily used aspects of it.

The AHP deals with problems of the following type.

A firm wishes to buy one new piece of equipment of a certain type and has four aspects in mind which will govern its purchasing choice: expense, E; operability, O; reliability, R; and adaptability for other uses, or flexibility, F. Competing manufacturers of that equipment have offered three options, X, Y and Z. The firm's engineers have looked at these options and decided that X is cheap and easy to operate but is not very reliable and could not easily be adapted to other uses. Y is somewhat more expensive, is reasonably easy to operate, is very reliable but not very adaptable. Finally, Z is very expensive, not easy to operate, is a little less reliable than Y but is claimed by the manufacturer to have a wide range of alternative uses. (This is obviously a hypothetical example and, to understand Saaty properly, you should think of another case from your own experience.)

Each of X, Y and Z will satisfy the firm's requirements to differing extents so which, overall, best meets this firm's needs?

This is clearly an important and common class of problem and the AHP has numerous applications but also some limitations which will be discussed at the end of this section.

Before giving some worked examples of the AHP, we need first to explain the underlying ideas. You do not need to understand matrix algebra to follow the line of argument but you will need that mathematical ability actually to apply the AHP. Take heart, this is the only part of the book which uses any mathematics.

THE BASIC PRINCIPLES OF THE AHP

The mathematics of the AHP and the calculation techniques are briefly explained in Annex A but its essence is to construct a matrix expressing the relative values of a set of attributes. For example, what is the relative importance to the management of this firm of the cost of equipment as opposed to its ease of operation? They are asked to choose whether cost is very much more important, rather more important, as important, and so on down to very much less important, than operability. Each of these judgements is assigned a number on a scale. One common scale (adapted from Saaty) is:

Intensity of importance	Definition	Explanation
1	Equal importance	Two factors contribute equally to the objective
3	Somewhat more important	Experience and judgement slightly favour one over the other.
5	Much more important	Experience and judgement strongly favour one over the other.
7	Very much more important	Experience and judgement very strongly favour one over the other. Its importance is demonstrated in practice.
9	Absolutely more important.	The evidence favouring one over the other is of the highest possible validity.
2,4,6,8	Intermediate values	When compromise is needed

Table 1 The Saaty Rating Scale

A basic, but very reasonable, assumption is that if attribute A is absolutely more important than attribute B and is rated at 9, then B must be absolutely less important than A and is valued at 1/9.

These pairwise comparisons are carried out for all factors to be considered, usually not more than 7, and the matrix is completed. The matrix is of a very particular form which neatly supports the calculations which then ensue (Saaty was a *very* distinguished mathematician).

The next step is the calculation of a list of the relative weights, importance, or value, of the factors, such as cost and operability, which are relevant to the problem in question (technically, this list is called an eigenvector). If, perhaps, cost is very much more important than operability, then, on a simple interpretation, the cheap equipment is called for though, as we shall see, matters are not so straightforward. The final stage is to calculate a Consistency Ratio (CR) to measure how consistent the judgements have been relative to large samples of purely random judgements. If the CR is much in excess of 0.1 the judgements are untrustworthy because they are too close for comfort to randomness and the exercise is valueless or must be repeated. It is easy to make a minimum number of judgements after which the rest can be calculated to enforce a perhaps unrealistically perfect consistency.

The AHP is sometimes sadly misused and the analysis stops with the calculation of the eigenvector from the pairwise comparisons of relative importance (sometimes without even computing the CR!) but the AHP's true subtlety lies in the fact that it is, as its name says, a *Hierarchy* process. The first eigenvector has given the relative importance attached to requirements, such as cost and reliability, but different machines contribute to differing extents to the satisfaction of those requirements. Thus, subsequent matrices can be developed to show how X, Y and Z respectively satisfy the needs of the firm. (The matrices from this lower level in the hierarchy will each have their own

eigenvectors and CRs.) The final step is to use standard matrix calculations to produce an overall vector giving the answer we seek, namely the relative merits of X, Y and Z vis-à-vis the firm's requirements.

A WORKED EXAMPLE

We know from the Introduction to this section that the firm has four factors in mind: expense, operability, reliability and flexibility; E, O, R and F respectively. The factors chosen should be independent, as required by Saaty's mathematics. At first sight, E and R are not independent but, in fact, what is really shown is that the firm would prefer not to spend too much money but is willing to do so if the results justify it.

We first provide an initial matrix for the firm's pairwise comparisons in which the principal diagonal contains entries of 1, as each factor is as important as itself.

	E	O	R	F
E	1			
O		1		
R			1	
F				1

There is no standard way to make the pairwise comparison but let us suppose that the firm decides that O, operability, is slightly more important than cost. In the next matrix that is rated as 3 in the cell O,E and 1/3 in E,O. They also decide that cost is far more important than reliability, giving 5 in E,R and 1/5 in R,E, as shown below.

	E	O	R	F
E	1	1/3	5	
O	3	1		
R	1/5		1	
F				1

The firm similarly judges that operability, O, is much more important than flexibility, F (rating = 5), and the same judgement is made as to the relative importance of F vis-à-vis R. This forms the completed matrix, which we will term the Overall Preference Matrix (OPM), is:

	E	O	R	F
E	1	1/3	5	1
O	3	1	5	1
R	1/5	1/5	1	1/5
F	1	1	5	1

The eigenvector (a column vector but written as a row to save space), which we will call the Relative Value Vector (RVV), is calculated by standard methods (see Annex A) as (0.232, 0.402, 0.061, 0.305). These four numbers correspond, in turn, to the relative values of E, O, R and F. The 0.402 means that the firm values operability most of all; 0.305 shows that they like the idea of flexibility; the remaining two numbers show that

they not desperately worried about cost and are not interested in reliability. The CR is 0.055, well below the critical limit of 0.1, so they are consistent in their choices. It may seem odd not to be interested in reliability but the RVV captures all the implicit factors in the decision context. Perhaps, in this case, the machine will only be used occasionally so there will be plenty of time for repairs if they are needed.

We now turn to the three potential machines, X, Y and Z. We now need four sets of pairwise comparisons but this time in terms of how well X, Y and Z perform in terms of the four criteria, E, O, R and F.

The first table is with respect to E, expense, and ranks the three machines as :

	X	Y	Z
X	1	5	9
Y	1/5	1	3
Z	1/9	1/3	1

This means that X is considerably better than Y in terms of cost and even more so for Z. Actual cost figures could be used but that would distort this matrix relative to others in which qualitative factors are assessed. The eigenvector for this matrix is (0.751, 0.178, 0.071), very much as expected, and the CR is 0.072, so the judgements are acceptably consistent.

The next three matrices are respectively judgments of the relative merits of X, Y and Z with respect to operability, reliability and flexibility (just to remind you, X is cheap and easy to operate but is not very reliable and could not easily be adapted to other uses. Y is somewhat more expensive, is reasonably easy to operate, is very reliable but not very adaptable. Finally, Z is very expensive, not easy to operate, is a little less reliable than Y but is claimed to have a wide range of alternative uses):

Operability:

	X	Y	Z
X	1	1	5
Y	1	1	3
Z	1/5	1/3	1

Eigenvector (0.480, 0.406, 0.114), CR=0.026

Reliability:

	X	Y	Z
X	1	1/3	1/9
Y	3	1	1/3
Z	9	3	1

Eigenvector (0.077, 0.231, 0.692), CR=0 (perfect consistency)

Flexibility:

	X	Y	Z
X	1	1/9	1/5
Y	9	1	2
Z	5	1/2	1

Eigenvector (0.066, 0.615, 0.319), CR= 0.

The reason that Y scores better than Z on this criterion is that the firm does not really believe the manufacturer's claims for Z. The AHP deals with opinion and hunch as easily as with fact.

The final stage is to construct a matrix of the eigenvectors for X, Y and Z

	E	O	R	F
X	0.751	0.480	0.077	0.066
Y	0.178	0.406	0.231	0.615
Z	0.071	0.114	0.692	0.319

This matrix, which we call the Option Performance Matrix (OPM), summarises the respective capability of the three machines in terms of what the firm wants. Reading down each column, it somewhat states the obvious: X is far better than Y and Z in terms of cost; it is a little better than Y in terms of operability, however, X is of limited value in terms of reliability and flexibility. These are not, however, absolutes; *they relate only to the set of criteria chosen by this hypothetical firm*. For another firm to whom reliability was more important and who wanted to avoid expense, the three machines might score quite differently.

Those results are only part of the story and the final step is to take into account the firm's judgements as to the relative importance of E, O, R and F. For a firm whose only requirement was for flexibility, Y would be ideal. Someone who valued only reliability would need machine Z. This firm is, however, more sophisticated, as, I suspect, are most firms, and has already expressed its assessment of the relative weights attached to E, O, R and F in the Relative Value Vector (0.232, 0.402, 0.061, 0.305). Finally, then, we need to weight the value of achieving something, R, say, by the respective abilities of X, Y and Z to achieve R, that is to combine the Relative Value Vector (RVV) with the Option Performance Matrix (OPM). Technically, the calculation is to post-multiply the OPM by the RVV to obtain the vector for the respective abilities of these machines to meet the firm's needs. It comes out to (0.392, 0.406, 0.204) and might be called the Value For Money vector (VFM). In matrix algebra, $OPM * RVV = VFM$ or, in words,

performance*requirement= value for money.

In those terms, this suggested method might have wide applicability.

The three numbers in the VFM are the final result of the calculation, but what do they mean?

First, in simple terms, they mean that X, which scores 0.392, seems to come out slightly worse in terms of its ability to meet the firm's needs than does Y at 0.406. Z is well behind at 0.204 and would do rather badly at satisfying the firm's requirements in this illustrative case.

Secondly, the three decimal places are, in practical terms, illusory, and X and Y are equal at 0.4. A coin could be tossed but, in the real world, it might be sensible to go for X as the option putting least pressure on cash flow.

Thirdly, and perhaps most importantly, the vector of the relative merits of X, Y and Z follows ineluctably from judgements made by the firm as to its requirements and by their engineers as to the capabilities of differing machines. There is a strong audit trail from output back to inputs. Of course, anyone who understands the AHP mathematics might be able to fiddle the judgements so as to guarantee a preferred outcome, but that is unavoidable expect by vigilance and the Delphi approach discussed below.

A SECOND EXAMPLE

Let us now look at the AHP in a different light.

Another firm has a different set of objectives. In their view, E is more important than O, but R and F are respectively much more important and absolutely more important than expense. They also judge that O is more important than R, that flexibility is more important than operability and that reliability is more important than flexibility. That all sounds rather confused and the AHP will help us to see just how confused it is.

The Overall Preference Matrix is:

	E	O	R	F
E	1	3	1/5	1/9
O	1/3	1	3	1/3
R	5	1/3	1	3
F	9	3	1/3	1

The eigenvector, or Relative Value Vector, turns out to be (0.113, 0.169, 0.332, 0.395) but the Consistency Index is 0.94, vastly in excess of the cut-off of 0.1 and indicating that the firm's valuations have, for all practical purposes, been made at random. (This example illustrates the immense importance of the calculation of the CR, a step which is sometimes omitted in careless use of the AHP.)

A set of preferences such as these are a recipe for a poor choice of machine and for endless "I told you so" afterwards. The explanation above foreshadows that but the calculation has confirmed just how incoherent the objectives are, and has done so in a way which might not have been so clear by verbal discussion. All too often, this sort of confusion remains hidden in the mind of the firm or, and even worse, in the separate minds of different interest groups within the organisation.

What is to be done in such a case?

The firm still needs a machine so the sensible course is to try to work out a consistent set of objectives. That could be supported by simple AHP software which draws attention to inconsistent choices enabling one to say “we’re not quite consistent yet”.

Let us suppose that rethinking the objectives leads to the new matrix:

	E	U	R	P
E	1	1/3	1/9	1/5
U	3	1	1	1
R	9	1	1	3
P	5	1	1/3	1

(It may help to understand the method for readers to work out for themselves the judgements to which these numbers correspond. It is meaningless to debate whether or not these are ‘good’ judgements; they reflect the firm’s mental model of the significance of the problem to the wider objectives of the firm).

Calculations from this matrix produce the Relative Value Vector, or eigenvector, (0.262, 0.454, 0.226) and a CR of 0.06, well on the safe side of the cut-off of 0.1; the judgements are now strongly consistent as opposed to the first set which were virtually random.

Weighting this RVV by the OPM previously calculated gives a Value for Money vector of the relative merits of machines X, Y and Z of (0.220, 0.360, 0.420). X is now well out of the running and Z is rather better than Y, but not dominantly so.

JUDGEMENTS IN THE AHP

The four factors used here, E, O, R and F were, of course, purely to demonstrate a calculation, but how might factors be determined in a real case? They could be an *ex cathedra* statement from someone in authority, but a more rational approach might be discussion with a small group, first in Focus Group mode to identify factors and then as a simple Delphi to obtain the Overall Preference Matrix. Recall from Chapter 3 that Delphi is a controlled debate and the reasons for extreme values are debated, not to force consensus, but to improve understanding.

STRENGTHS AND WEAKNESSES OF THE AHP

Like all modelling methods, the AHP has strengths and weaknesses.

The main advantage of the AHP is its ability to rank choices in the order of their effectiveness in meeting conflicting objectives. If the judgements made about the relative importance of, in this example, the objectives of expense, operability, reliability and flexibility, and those about the competing machines’ ability to satisfy those objectives, have been made in good faith, then the AHP calculations lead inexorably to

the logical consequence of those judgements. It is quite hard – but not impossible – to ‘fiddle’ the judgements to get some predetermined result. (In MOA, it *is* impossible to do that.) The further strength of the AHP is its ability to detect inconsistent judgements.

The limitations of the AHP are that it only works because the matrices are all of the same mathematical form – known as a positive reciprocal matrix. The reasons for this are explained in Saaty’s book, which is not for the mathematically daunted, so we will simply state that point. To create such a matrix requires that, if we use the number 9 to represent ‘A is absolutely more important than B’, then we have to use 1/9 to define the relative importance of B with respect to A. Some people regard that as reasonable; others are less happy about it.

The other seeming drawback is, that if the scale is changed from 1 to 9 to, say, 1 to 29, the numbers in the end result, which we called the Value For Money Vector, will also change. In many ways, that does not matter as the VFM (not to be confused with the Viable Final Matrix) simply says that something is relatively better than another at meeting some objective. In the first example, the VFM was (0.392, 0.406, 0.204) but that only means that machines A and B are about equally good at 0.4, while C is worse at 0.2. It does not mean that A and B are twice as good as C.

In less clear-cut cases, it would be no bad thing to change the rating scale and see what difference it makes. If one option consistently scores well with different scales, it is likely to be a very robust choice.

In short, the AHP is a useful technique for discriminating between competing options in the light of a range of objectives to be met. The calculations are not complex and, while the AHP relies on what might be seen as a mathematical trick, you don’t need to understand the maths to use the technique. Do, though, be aware that it only shows *relative* value for money.

NOTES AND REFERENCES

Coyle R G (1989) 'A Mission-orientated Approach to Defense Planning', *Defense Planning*, Vol. 5, No.4. pp 353-367.

Miller, D. (1998) *The Cold War: A Military History*, London, John Murray.

Saaty T L (1980) *The Analytic Hierarchy Process*, McGraw Hill International.

ANNEX A

THE AHP THEORY AND CALCULATIONS

The mathematical basis of the AHP can be explained in fairly simple outline for the purposes of this book but you need to know what a matrix and a vector are and how to multiply a matrix by a vector. For a full treatment of the AHP the mathematically undaunted should refer to Saaty's book. We will cover the mathematics first and then explain the calculations.

THE AHP THEORY

Consider n elements to be compared, $C_1 \dots C_n$ and denote the relative 'weight' (or priority or significance) of C_i with respect to C_j by a_{ij} and form a square matrix $A=(a_{ij})$ of order n with the constraints that $a_{ij} = 1/a_{ji}$, for $i \neq j$, and $a_{ii} = 1$, all i . Such a matrix is said to be a reciprocal matrix.

The weights are consistent if they are transitive, that is $a_{ik} = a_{ij}a_{jk}$ for all i, j , and k . Such a matrix might exist if the a_{ij} are calculated from exactly measured data. Then find a vector ω of order n such that $A\omega = \lambda\omega$. For such a matrix, ω is said to be an eigenvector (of order n) and λ is an eigenvalue. For a consistent matrix, $\lambda = n$.

For matrices involving human judgement, the condition $a_{ik} = a_{ij}a_{jk}$ does not hold as human judgements are inconsistent to a greater or lesser degree. In such a case the ω vector satisfies the equation $A\omega = \lambda_{\max}\omega$ and $\lambda_{\max} \geq n$. The difference, if any, between λ_{\max} and n is an indication of the inconsistency of the judgements. If $\lambda_{\max} = n$ then the judgements have turned out to be consistent. Finally, a Consistency Index can be calculated from $(\lambda_{\max}-n)/(n-1)$. That needs to be assessed against judgments made completely at random and Saaty has calculated large samples of random matrices of increasing order and the Consistency Indices of those matrices. A true Consistency Ratio is calculated by dividing the Consistency Index for the set of judgments by the Index for the corresponding random matrix. Saaty suggests that if that ratio exceeds 0.1 the set of judgments may be too inconsistent to be reliable. In practice, CRs of more than 0.1 sometimes have to be accepted. A CR of 0 means that the judgements are perfectly consistent.

THE AHP CALCULATIONS

There are several methods for calculating the eigenvector. Multiplying together the entries in each row of the matrix and then taking the n^{th} root of that product gives a very good approximation to the correct answer. The n^{th} roots are summed and that sum is used to normalise the eigenvector elements to add to 1.00. In the matrix below, the 4th root for the first row is 0.293 and that is divided by 5.024 to give 0.058 as the first element in the eigenvector.

The table below gives a worked example in terms of four attributes to be compared which, for simplicity, we refer to as A, B, C, and D.

	A	B	C	D	n th root of product of values	Eigenvector
A	1	1/3	1/9	1/5	0.293	0.058
B	3	1	1	1	1.316	0.262
C	9	1	1	3	2.279	0.454
D	5	1	1/3	1	1.136	0.226
Totals					5.024	1.000

The eigenvector of the relative importance or value of A, B, C and D is (0.058,0.262,0.454,0.226). Thus, C is the most valuable, B and D are behind, but roughly equal and A is very much less significant.

The next stage is to calculate λ_{\max} so as to lead to the Consistency Index and the Consistency Ratio.

We first multiply on the right the matrix of judgements by the eigenvector, obtaining a new vector. The calculation for the first row in the matrix is:

$$1*0.058+1/3*0.262+1/9*0.454+1/5*0.226 = 0.240$$

and the remaining three rows give 1.116, 1.916 and 0.928. This vector of four elements (0.240,1.116,1.916,0.928) is, of course, the product $A\omega$ and the AHP theory says that $A\omega=\lambda_{\max}\omega$ so we can now get four estimates of λ_{\max} by the simple expedient of dividing each component of (0.240,1.116,1.916,0.928) by the corresponding eigenvector element. This gives $0.240/0.058=4.137$ together with 4.259, 4.22 and 4.11. The mean of these values is 4.18 and that is our estimate for λ_{\max} . If any of the estimates for λ_{\max} turns out to be less than n, or 4 in this case, there has been an error in the calculation, which is a useful sanity check.

The Consistency Index for a matrix is calculated from $(\lambda_{\max}-n)/(n-1)$ and, since $n=4$ for this matrix, the CI is 0.060. The final step is to calculate the Consistency Ratio for this set of judgements using the CI for the corresponding value from large samples of matrices of purely random judgments using the table below, derived from Saaty's book, in which the upper row is the order of the random matrix, and the lower is the corresponding index of consistency for random judgements.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0.00	0.00	0.58	0.90	1.12	1.24	1.32	1.41	1.45	1.49	1.51	1.48	1.56	1.57	1.59

For this example, that gives $0.060/0.90=0.0677$. Saaty argues that a $CR > 0.1$ indicates that the judgements are at the limit of consistency though $CRs > 0.1$ (but not too much more) have to be accepted sometimes. In this instance, we are on safe ground.

A CR as high as, say, 0.9 would mean that the pairwise judgements are just about random and are completely untrustworthy.