The essential multiobjectivity of linear programming

TJ Stewart*

Dedication to Emeritus Professor Gerhard Geldenhuys

In 1971 I joined the National Research Institute of Mathematical Sciences (NRIMS) of the CSIR, as the sole employee in an effectively one-person operations research section. This was, for me, a substantial career shift from chemical engineering into OR (a shift made by many at that time), but one which has been my life ever since. A useful source document available to me at that time in NRIMS, was the set of notes prepared a few years years earlier by Gerhard Geldenhuys [2] to "give a brief and elementary introduction to linear programming, game theory, queueing theory and dynamic programming". This document had introduced OR to NRIMS, and was a major stimulus to the creation of the OR section at NRIMS/CSIR, the forerunner of the current OR group at CSIR, and thus also to my career shift. In this sense, Gerhard Geldenhuys can be seen as the father of a centre of OR which has arguably been one of the most influential in this country over the past four decades.

Abstract

It is argued that any non-trivial real world problems involve multiple objectives. The simplistic approach of combining objectives in linear form can generate highly misleading and biased results, and is poor operational research practice. Such biases are illustrated by means of a simple example, and it is demonstrated that two simple approaches based on goal programming and piecewise linear value functions can largely resolve these problems. These methods require careful attention to problem structuring but are easily implemented by use of standard linear programming (LP) packages.

1 Background

Two of the four themes in [2] are the closely related issues of linear programming (LP) and game theory. Many textbooks on OR/MS put such an emphasis on LP that students

---

*(**Fellow of the Operations Research Society of South Africa**) Department of Statistical Sciences, University of Cape Town, Private Bag, Rondebosch 7701, South Africa, email: `Theodor.Stewart@uct.ac.za`

may be forgiven for viewing OR and LP as almost synonymous. While such a view of OR represents a dangerously limited and restricted perspective of the subject, there is no doubt that LP remains one of the workhorses in the practice of OR. The basic thesis of the present paper, however, is that LP is all too often poorly and ineffectively used, largely because of a lack of recognition that a single objective is the exception rather than the rule in all but the most trivial OR applications.

Standard textbooks in OR start with some reference to an "OR process" that includes identification of alternatives and objectives, and comparing alternatives in terms of these objectives. For example, Winston [13] recognizes a "seven-step model-building process" consisting of (1) Formulate the problem; (2) Observe the system; (3) Formulate a mathematical model of the problem (although this step may be challenged by those from the "soft OR" school); (4) Verify the model and use the model for prediction; (5) Select a suitable alternative (in which Winston explicitly recognizes that there *may* be more than one objective); (6) Present the results and conclusions of the study to the organization; and (7) Implement and evaluate recommendations.

Such an OR process inevitably acknowledges careful problem structuring as an essential part of OR (*e.g.* Winston's steps 1–3), and to the recognition of "objectives" in the plural. Nevertheless, most texts then launch into linear programming with a strong emphasis on algorithmic methods (and who, today, needs to know much about the tricks of the classical simplex algorithm?), with an almost default assumption that the objective function is well-defined and self-evident. Any admission of the existence of multiple objectives will typically be relegated to fairly superficial mention in a couple of sections on goal programming, often limited to discussion of preemptive goal programming (dating from the 1960s) and associated variations to the simplex method to provide a computationally efficient method of solution. The only other mention of multiple objectives in such texts may be a brief mention of the analytic hierarchy process (AHP) in a later chapter on decision theory, perhaps as an "example" of the use of eigenvalue theory!

Thus Winston [13] devotes five chapters to LP, but the only mention of dealing with multiple objectives is a brief final section in Chapter 4 on goal programming, essentially only dealing with the preemptive variant. The well-known text by Hillier and Lieberman [3] places a similarly brief mention of preemptive goal programming only at the end of Chapter 7. The otherwise excellent book of Ragsdale [7] does provide a more thorough treatment of goal programming and related issues, but this, too, only appears towards the end of Chapter 7. In other words, the treatment of multiple objectives in LP is dealt with very much as an afterthought, a special case that may occasionally arise. We shall argue that precisely the opposite is true.

Even in private sector applications with an unquestioning belief in the power of markets, however, a single financial objective is virtually never relevant as has been so clearly argued by Kaplan and Norton [5] in presenting their concept of the "balanced scorecard." Public sector applications (especially in a developing world context) inevitably have to deal with conflicts between economic, social and environmental goals as well as between the interests of different stakeholder groups. These multiple goals are seldom directly commensurate, and as we shall illustrate cannot generally be aggregated in simple additive fashion. We shall argue, therefore, that the standard for applying LP in practice *must* be to adopt a

multiobjective linear programming (MOLP) approach, and to apply one or more of the available methodologies of MOLP rather than some arbitrary unitary objective function even when the latter is subject to comprehensive sensitivity analysis.

In the next section, we illustrate by means of a simple example the manner in which a simple linear aggregate objective function leads to extreme solutions, which are seldom if ever acceptable to decision makers, leading often to an *ad hoc* introduction of additional constraints which address symptoms rather than causes of the problem. In Section 3 we present fundamental requirements that need to be incorporated into any (multiobjective) LP model in order to address the real problem. Two different approaches to satisfying these requirements, making use of (generalized) goal programming and piecewise linear value functions, are described in Sections 4 and 5. Section 6 demonstrates that these approaches are easily implemented in standard software, *i.e.* that there is no need for specialized software.

2 Linear aggregation of component objectives

The following problem was introduced in Chapter 3 of Belton and Stewart [1], and was a trimmed down version of the game reserve problem discussed earlier by Jordi and Peddie [4]. The decision variables are numbers (X_{sc}, expressed in 100s) of animals of species s to be stocked in each camp c, and numbers of excess stock (Y_{sj}) of species s to be disposed of by each of three available methods j (*e.g.* sales, hunting or culling for food). For purposes of numerical illustration, we suppose that there are only two camps, and that we are only considering three different species. Details are given in [1], but the constraints (for the trimmed down problem) are given by

$$\left.\begin{array}{rcl}
Y_{11} + Y_{12} + Y_{13} &=& 0.4X_{11} + 0.4X_{12} \\
Y_{21} + Y_{22} + Y_{23} &=& 0.3X_{21} + 0.3X_{22} \\
Y_{31} + Y_{32} + Y_{33} &=& 0.2X_{31} + 0.2X_{32} \\
X_{11} + 3X_{21} + 8X_{31} &\leq& 14 \\
2X_{11} + X_{21} + 4X_{31} &\leq& 12 \\
X_{12} + 3X_{22} + 8X_{32} &\leq& 25 \\
2X_{12} + X_{22} + 4X_{32} &\leq& 30 \\
X_{31} + X_{32} &\geq& 3 \\
X_{sc}, Y_{sj} &\geq& 0 \qquad \forall\ s, c, j.
\end{array}\right\}
\tag{1}$$

The first three constraints balance disposal of excess stocks with natural population growth. The next four constraints indicate maximum availability of two food stocks in each of two camps. The final constraint (apart from non-negativity) places a lower bound on numbers of one species that need to be kept for conservation reasons.

Standard textbook LP (*e.g.* those cited in the previous section) would now move towards formulating *the* objective function. For example, Hillier and Lieberman [3, p. 33] speak of constructing an overall measure of performance which is linear in the variables. In practice, however, such a measure of performance would need to include contributions from a number of different attributes such as (in the above example) revenues, contribution to conservation of endangered species, maintenance of good relations with local communities

through provision of food and/or traditional hunting rights, and good game viewing for tourists. For purposes of illustration, suppose that four attributes (or objectives) can be identified as follows, derived from the above-mentioned aspects:

$$\left.\begin{aligned}
z_1 &= 2Y_{12} + 2Y_{13} + 3Y_{22} + 6Y_{23} + Y_{32} + 8Y_{33} \\
z_2 &= 5Y_{12} + 3Y_{22} + 2Y_{32} \\
z_3 &= 20Y_{11} + 15Y_{21} \\
z_4 &= 3X_{11} + 15X_{21} + 6X_{31} + X_{12} + 5X_{22} + 2X_{32}.
\end{aligned}\right\} \quad (2)$$

How might the analyst then go about formulating the aggregate measure of performance, or objective? The quoted texts (and many others) are silent in this regard. However, the contribution of each decision variable to the overall objective would need (in order to preserve linearity) to be a linear sum of the contributions to each of the component output measures. In other words, the aggregate objective would need to be expressible in the form $w_1 z_1 + w_2 z_2 + w_3 z_3 + w_4 z_4$. Such an aggregation is generally not recognized explicitly in the textbooks, so that no guidance in choosing weights is given. The MCDM literature contains much discussion on the process of identifying importance weights, for which a good starting point is to establish the available ranges of options on each objective. Useful approximations to these ranges are often provided by the so-called "payoff" table, which is obtained by maximizing for each objective i an aggregate objective having $w_i = 1$ and $w_j = \epsilon$ for $j \neq i$ (for some small $\epsilon > 0$). The resulting payoff table for the above example is illustrated by Table 1, giving the values for each of the four objectives when effectively maximizing each in turn.

Attribute being maximized	Values obtained for:			
	z_1	z_2	z_3	z_4
z_1	16.80	30.00	0	33.00
z_2	12.60	31.20	0	33.00
z_3	4.80	0	120.00	33.00
z_4	4.80	0	22.50	77.67

Table 1:   *Payoff table for the game reserve example.*

In order to ensure comparability, the z_i might be rescaled to a consistent range (*e.g.* 0–1) within the payoff table, before applying importance weights. (For example, the first attribute may be redefined as $z_1' = (z_1 - 4.8)/(16.8 - 4.8)$.) For illustration, suppose that importance weights of 10, 8, 6 and 6 are allocated to the rescaled attributes. The composite weight for z_1, including the appropriate scaling constant for comparability of ranges, would then be 10/12=0.833. Extending this argument to all four objectives generates an objective function defined on the original attributes of $0.833z_1 + 0.256z_2 + 0.05z_3 + 0.134z_4$.

Up to this point, the analysis appears to have been quite careful and thoughtful, but the resulting solution gives $z_1 = 16.8$, $z_2 = 30$, $z_3 = 0$ and $z_4 = 33$. Comparison with Table 1 reveals how extreme this solution is — it is simply the solution found when maximizing z_1 in isolation, and is very good on two of the attributes and very bad on the other two. Sensitivity analysis on w_3, say, shows that the solution is unchanged as w_3 is increased to

virtually three times its original value. At this point the optimal solution jumps suddenly to another extreme, namely $z_1 = 4.8$, $z_2 = 0$, $z_3 = 120$ and $z_4 = 33$. No intermediate solution is ever discovered. Our experience is that this type of phenomenon is not a freak circumstance — it is almost inevitable with LP, whether the underlying components of the objective have been explicitly recognized or not.

What happens at this stage? Our experience suggests that when decision makers or analysts examine an extreme solution such as that illustrated in the previous paragraph, there is a tendency to respond by imposing additional constraints in order to make the solution "more realistic." If the key attributes as illustrated above are made explicit, the added constraints may be expressed directly in terms of these attributes (which is at least a goal-directed response). Often, however, such attributes are not made explicit, and constraints are applied more directly to decision variables. In the above example, the decision maker may observe that almost all of the disposal of excess stocks is by method 2 (which corresponded to live sales of animals), and may react by placing constraints on the numbers of live sales for each species.

The trouble with this approach (of adding constraints whenever the solution is perceived to be too extreme in some sense) is that it is entirely *ad hoc* (addressing symptoms rather than causes), and tends to need to be applied incrementally as other extreme phenomena emerge after the first has been corrected. There is a real danger that eventually the added constraints do no more than to force the solution into a preconceived form, providing a "rational" justification from the model for such preconception.

The core reason why extreme solutions are generally unsatisfactory is that the existence of multiple objectives is fundamental, and not just an occasional nuisance phenomenon. The effect of multiple objectives may be understood in different ways, although these are perhaps merely different sides of the same coin. For example, the following two explanations may often be useful in trying to understand the fundamental problem:

- *Decreasing marginal returns to scale:* The more that is gained on any one criterion, the less is the marginal value of further increases.

- *Simon's [9] satisficing principle:* Decision makers strive for improvement in any one criterion up to a point of "sufficient satisfaction," after which attention is turned to other criteria.

3 Fundamental modelling requirements

The primary themes of the previous section were:

1. that optimization models in OR should start from the basis that multiple objectives are the norm and not the exception; and

2. that the treatment of multiple objectives needs to recognize the principles of changing marginal returns to scale and/or satisficing in modelling preferences (as opposed to constant tradeoffs).

At an operational level, the implication for linear programming models arising from the above points is that a larger number of feasible bases needs to be created, in order to provide more opportunity for a richer exploration of less extreme solutions within the standard LP framework. The objective function must be constructed in such a way that the opportunities provided by such an enlarged set of feasible bases may be exploited to find the most satisfactory solution.

The simple expedient of dealing with multiple objectives, even when these have been explicitly defined, by a simple linear aggregation fails to make any contribution to the above modelling needs. This failure is clearly illustrated by the game reserve planning example of the previous section. Use of a weighted sum of objectives cannot be classified as *multiple criteria decision analysis (MCDA)* at all, and must be viewed as highly inadequate and poor OR practice.

A superficial reaction to the assertions of the previous two paragraphs may be that these are all well and good in theory, but that practitioners do not have the time for more complex models, and that in any case appropriate software may not easily be available. In the next two sections, we shall review two approaches which do not involve substantial complexities, and which are easily implemented in standard LP packages. What these approaches do demand is for additional attention to be paid to structuring of decision maker preferences, but that is surely fundamental to good, professional practice in OR.

Neither of the two approaches to be discussed introduce any fundamentally new concepts, but our aim is to examine the methods of implementation which best match up with the modelling requirements presented at the start of this section. Goal programming and the related methods discussed in Section 4 model preferences in terms of satisficing requirements, while the goal achievement constraints create the increased number of basic solutions. In fact, the Chebychev form of goal programming (or the Wierzbicki generalization thereof) creates an even richer range of basic solutions to explore.

The piecewise linear value function approach of Section 5 models preferences directly in terms of changing tradeoffs as relative levels of achievement change. The desired enrichment of the set of feasible basic solutions is realized by the constraints needed to define the different segments which approximate a non-linear value function.

4   Goal programming and extensions

The original goal programming (GP) model was an early extension of LP, and is generally discussed in introductory texts on OR/MS, so that most readers would probably consider this approach to be very well known. Nevertheless, there are subtleties that seem often to be overlooked, especially as many presentations of GP focus on the algorithmic features viewed as an extension of linear programming, rather than on the preference modelling aspects which are much more fundamental to the practice of OR.

GP is frequently seen as modelling "soft" or imprecise constraints, for example that return on investment should be about 10%, say. This use of GP is an occasionally useful device, but the real power of GP is in recognizing and modelling conflicting decision objectives. So let us suppose that we have recognized m measures of performance that can be represented

in linear form by

$$\sum_{j=1}^{n} c_{ij}x_j, \qquad i = 1, \ldots, m, \tag{3}$$

where the x_j are the decision variables (required to satisfy some set of linear constraints).

A common introduction to goal programming would define goal levels of performance, say g_i for $i = 1, \ldots, m$, and include constraints of the form

$$\sum_{j=1}^{n} c_{ij}x_j + \delta_i^- - \delta_i^+ = g_i, \qquad i = 1, \ldots, m, \tag{4}$$

where the $\delta_i^- > 0$ and $\delta_i^+ > 0$ represent deviations below and above the goal (target) respectively. Although the associated discussion might well emphasize that in many cases deviations to one side of g_i may be substantially more important than to the other side, the impression conveyed is still that of the goal being an ideal to be striven for.

In structuring decision problems in multiple criteria terms, however, it is more usual that the measures of performance defined by (3) are associated with a *direction* of preference. Without loss of generality we may assume that the coefficients are defined such that increasing values of the measures of performance are always desirable. The role of the goal value g_i is to indicate the level of performance beyond which other criteria are likely to be more important to final decision making. It would be more natural, and more transparent to the decision maker, then to replace (4) by

$$\sum_{j=1}^{n} c_{ij}x_j + \delta_i \geq g_i, \qquad i = 1, \ldots, m, \tag{5}$$

where $\delta_i > 0$ now clearly represents underachievement on the associated goal.

In fact, we have argued elsewhere (Belton and Stewart [1], pp. 210–211) that formulations in terms of a central desirable target (*e.g.* preferred temperature of your coffee) can generally be decomposed into two objectives working in opposite directions, and that explicit separation of these two objectives is better decision modelling practice. Suppose, for example, that one measure of performance relates to a quality measure such as proportion of contaminants in a product. Suppose further that a goal of 5% is stated, and that deviations to either side are said to be less desirable. What does this signify at a more fundamental level? It probably means that the decision maker is concerned both about costs or throughput (if the proportion is made unnecessarily small) and about customer relations (if the proportion becomes too large). The target level then represents the decision maker's *a priori* judgement as to what level of compromise is likely to balance the two objectives satisfactorily. The best compromise should, however, emerge *from* the analysis rather than to be imposed *on* the analysis. With this aim in mind, it would be better to identify and to deal with both objectives separately (even if based on the same numerical measure of performance), rather than to aggregate too soon in the process.

For this reason, the remainder of the discussion will assume a clearly directional objective (maximization for purposes of discussion), and with goals and deviations defined through (5).

The core principle in goal programming is to minimize the magnitude of the deviations according to some appropriate norm. Many texts spend time discussing preemptive or lexicographic goal programming, in which the goals are ordered by importance, after which each deviation is minimized in turn, fixing the previous (more important) deviations at their minimal values. The early goal programming literature developed a quite clever extension to the simplex algorithm which enabled preemptive GP to be executed in essentially a single call to the modified simplex algorithm. This extension is hardly relevant today as few if any OR specialists attempt to implement their own simplex code. As a modelling strategy, preemptive GP assumes that there is no compensation between goal achievements for different objectives. No gain on a criterion of lesser importance can ever compensate for even an infinitesimal loss on a criterion at a higher importance level. It is, however, difficult to imagine contexts in which such a total lack of compensation between objectives is a realistic representation of decision maker preferences. (For example, we all gamble with our lives for the convenience of using a car to go the shops on a rainy day!)

The acceptance of compensation between objectives does require the specification of *weights* to be applied to the deviations, *i.e.* such that for any two goals, equally weighted deviations ($w_i \delta_i$, say) represent deviations of equal importance to the decision maker. The concepts of "*swing weighting*" from value function theory (*e.g.* Belton and Stewart [1], Section 5.4.2) do provide operationally practical means for establishing suitable weights. Fortunately, however, the results from goal programming models tend not to be overly sensitive to weights as long as the orders of magnitude are comparable, and we shall not dwell on the assessment of weights in this paper.

The simplest norm for assessing the aggregate magnitude of (weighted) goal deviations is the Archimedean, or simple weighted sum

$$\sum_{i=1}^{m} w_i \delta_i. \tag{6}$$

Minimization of (6) subject to the constraints on the x_j and to the goal achievement constraints (5) then provides the (Archimedean) GP solution by use of standard LP packages. In spite of the superficial similarities to the simple linear weighted aggregation which we earlier dismissed, results tend to be much more satisfying and less extreme. In fact, although we shall motivate for use of the Chebychev norm in the following paragraphs, it must be conceded that in many cases the two norms lead to very similar results.

Minimization of (6) can, however, in some instances lead to situations in which one deviational variable (δ_i) is relatively large, while all (or most) others are 0. This phenomenon might still produce results with less balance between goal achievements than is fully desired. There may then be advantage in penalizing large deviations relatively more severely than small deviations. Possibilities such as a Euclidean norm ($\sum_{i=1}^{m} [w_i \delta_i]^2$) might address the problem, but the shift to a non-linear objective creates other disadvantages in computational efficiency and in access to standard post-optimality analysis. An alternative norm which achieves the desired end without sacrificing linearity is the Chebychev norm

$$\max_{i=1}^{n} \{w_i \delta_i\}.$$

Minimization of the Chebychev norm is achieved by introducing one further variable (say Δ), and minimizing Δ subject to

$$\Delta \geq w_i \delta_i, \qquad i = 1, \ldots, m \tag{7}$$

and to all previously defined constraints.

It needs to be recognized, however, that even if for a particular solution the largest deviation cannot be improved upon, there may yet remain opportunities to improve achievements in other goals. Such opportunities may not be recognized by simply minimizing the maximum deviation, so that the resulting solution may be inefficient. In order to avoid such inefficiencies, it is usual to replace minimization of Δ by minimization of

$$\Delta + \epsilon \sum_{i=1}^{m} w_i \delta_i \tag{8}$$

for some suitably small $\epsilon > 0$. Minimization of (8) subject to the constraints on x_j and to (5) and (7) then generates the Chebychev GP solution, again requiring only the use of standard LP packages.

The process of finding the Chebychev GP solution focuses in effect at any one stage on improving the relatively worst weighted deviation. Such an approach does seem to be closer to the spirit of "satisficing" than the other GP approaches. Experience suggests that where the Archimedean and Chebychev GP solutions differ substantially, it is the Chebychev form which generates the more balanced solutions. For these reasons, the present author generally recommends using the Chebychev form as the standard implementation of GP. In essence, the approach introduces $m + 1$ additional variables and $2m$ additional constraints into the LP. As the number of objectives is typically orders of magnitude smaller than the numbers of decision variables (x_j) and the numbers of constraints on the variables, the increase in problem size is typically hardly detectable.

The Chebychev GP approach is closely linked to Wierzbicki's [12] generalization of goal programming, which he terms reference point methods. Wierzbicki recognized that in some cases it may be possible to achieve performance better than the predefined goals (which he terms the "reference point"). The same arguments used above to justify the Chebychev model in terms of focus on the worst performing objective can be extended beyond the reference point if that is possible. In effect, Wierzbicki's approach can be seen as precisely the Chebychev GP formulation, but with the removal of the non-negativity constraints on the δ_i and Δ. This too is easily implemented in standard LP packages.

An illustration of implementation in standard software is given in Section 6.

5 Piecewise linear value functions

Following the same problem structure as that used in the previous section, we now define an explicit performance measure variable for each objective as

$$z_i = \sum_{j=1}^{n} c_{ij} x_j \qquad i = 1, \ldots, m. \tag{9}$$

With the same assumptions as before, it is desired to maximize each z_i in some sense.

Standard axioms of value function theory (*e.g.* Belton and Stewart [1], Section 4.2, or Keeney and Raiffa [6], Chapter 3) imply that provided that decision maker preferences are complete and transitive and that the objectives are *preferentially independent*, there will exist an additive value function which will describe decision maker preferences. In other words, there exist functions $v_i(z_i)$ such that the decision maker will prefer an outcome described by $\mathbf{z}^a = (z_1^a, z_2^a, \ldots, z_m^a)$ to one described by the vector $\mathbf{z}^b$ if and only if $V(\mathbf{z}^a) > V(\mathbf{z}^b)$, where

$$V(\mathbf{z}) = \sum_{i=1}^{m} v_i(z_i). \tag{10}$$

We note that the expression of a value function in this form implies that importance of an objective is represented in the scaling of the associated partial value function $v_i(z_i)$. To make this point explicit for later reference, suppose that the range of possible outcomes to be considered for z_i has been determined (perhaps by means of a payoff table such as illustrated in Section 2) to be given by z_i^0 (minimum) to z_i^* (maximum). If the usual convention of setting $v_i(z_i^0) = 0$ is adopted, then $v_i(z_i^*)$ would be a measure of the importance of objective i.

In general, the marginal value functions may be expected to be non-linear, as a linear function would imply constant marginal values (tradeoffs with other objectives) at all levels of performance. As indicated earlier, constant marginal values would be an exception rather than the rule. Furthermore, as demonstrated in other work (*e.g.* Stewart [10, 11]), the inappropriate use of linear approximations can severely distort the results obtained. If decreasing marginal returns to scale do apply, however, we can conclude that the marginal value functions must be concave.

If the value functions could be established directly, then the maximum value solution would be obtained by maximizing (10) subject to the constraints. Two practical problems arise:

1. complete specification of the marginal value functions can be quite a time-consuming process; and

2. the resulting optimization is non-linear, losing much of the power of the LP algorithms.

Stewart [10, 11], however, demonstrated that the approximation of each $v_i(z_i)$ by a piecewise linear function with four linear segments provides very close to optimal solutions in most cases (certainly relative to the imprecisions and uncertainties in other aspects of the model specification). Piecewise linear functions are easily incorporated into the LP structure. For each objective, we define four new variables, say $z_{i1}, \ldots, z_{i4}$ with the restriction that

$$z_i = z_i^0 + \sum_{k=1}^{4} z_{ik}$$

or, more directly in terms of the original decision variables,

$$\sum_{k=1}^{4} z_{ik} = \sum_{j=1}^{n} c_{ij}x_j - z_i^0 \qquad i = 1, \dots, m. \tag{11}$$

The marginal value function is then approximated by

$$v_i(z_i) \approx \sum_{k=1}^{4} \beta_{ik} z_{ik}, \tag{12}$$

where β_{ik} is the approximated slope of $v_i(z_i)$ over the k'th segment.

If equal-lengthed segments are used (which is not necessary, but facilitates presentation here), then the portion of z_i allocated to each segment must be bounded above by $\lambda_i = (z_i^* - z_i^0)/4$. We should also ensure the precedence relationships in that $z_{ik} > 0$ only if $z_{i,k-1} = \lambda_i$ for $k > 1$. If the marginal value functions are concave (as suggested above they might well be), then the precedence relationships will automatically be satisfied at the optimal solution, since we would have $\beta_{i,k-1} > \beta_{ik}$ for $k > 1$. If more general marginal value functions arise (*e.g.* sigmoidal), then precedence will need to be ensured by introducing binary variables. For purposes of presentation here, we shall restrict ourselves to the concave case.

All that remains is to specify the parameters β_{ik}, which may be done relatively easily. We note firstly that $v_i(z_i^*)$ is approximated by $\sum_{k=1}^{4} \beta_{ik}\lambda_i$ which is intended as measure of relative importance of objective i. The relative importance of each objective can be assessed as described in both Sections 2 and 4 (by "swing weighting"). Suppose that importance weights (on an arbitrary scale) are assessed as $W_1, \dots, W_m$. Then the parameters need to be chosen so as to satisfy

$$\sum_{k=1}^{4} \beta_{ik} \propto W_i/\lambda_i.$$

Complete specification of the parameters is achieved by establishing the relative importances within each objective of the gains across each succeeding segment.

As an example, consider again the game reserve planning problem. We had suggested for illustrative purposes relative importances of 10:8:6:6, so that we could for the assessment above start with $W_1 = 10$, $W_2 = 8$ and $W_3 = W_4 = 6$. If the z_i^* and z_i^0 are chosen as the maximum and minimum in each column of Table 1, then we easily obtain $\lambda_1 = 3$, $\lambda_2 = 7.8$, $\lambda_3 = 30$ and $\lambda_4 = 11.17$.

Suppose further that the decision maker is satisfied with the assumption of decreasing marginal returns on each objective, so that $\beta_{i1} > \beta_{i2} > \beta_{i3} > \beta_{i4}$ for each i. Finally suppose that (perhaps by displaying alternative shapes of the function graphically), the decision maker accepts that all four functions have a similar shape, with the slopes over the four segments in the ratios 10:7:4:2. We could start by setting $\beta_{i1} = W_i/\lambda_i$ for each i. Then β_{i2} would need to be given by $0.7\beta_{i1}$, and so on for the other segments.

The piecewise linear approximation to the optimal solution is then obtained by maximizing $\sum_{i=1}^{m} \sum_{k=1}^{4} \beta_{ik} z_{ik}$, subject to the original constraints, to (11) and to the requirements that

	A	B	C	D	E	F	G	H	I	J	K
1		Numbers	(decision variables):				Food Consumptions:			Food Available	
2			Species:					Species:			
3		A	B	C			A	B	C	Camp1	Camp2
4	Camp1	4.00	2.00	0.50		Stock1	1.00	3.00	8.00	14.00	25.00
5	Camp2	5.00	0.00	2.50		Stock2	2.00	1.00	4.00	12.00	30.00
6	Meat	0.00	0.00	0.00						=>=	=>=
7	LiveSales	3.60	0.60	0.11						=>=	>=
8	Hunting	0.00	0.00	0.49		Total consumed:	Stock1			14.00	25.00
9							Stock2			12.00	20.00
10											
11	Species Totals	9.00	2.00	3.00							
12	Growth Rate	0.40	0.30	0.20							
13	Production	3.60	0.60	0.60		Species C Total		3.00	=>=	3.00	
14		=	=	=							
15	Disposal	3.60	0.60	0.60							
16											
17	Revenues (z1)	13.00									
18	Conservation (z2)	20.03									
19	Local Food (z3)	0.00									
20	Viewing (z4)	55.00									
21											

Figure 1: *Spreadsheet set-up for the game reserve planning LP.*

$z_{ik} \leq \lambda_i$ for all i and $k < 4$. At least for the concave function case, this is a standard LP structure. The extended structure has introduced $4m$ new variables, m new linear constraints, and $3m$ additional upper bounds on variables. As with goal programming, this represents a very modest expansion to the size of the problem.

6   Software implications

The question is often raised concerning the availability of software for multiple objective linear programming. The simple lesson from the previous two Sections is that there is no need for specialized software. Both the goal programming and piecewise linear value function approaches are easily implemented in standard LP packages. This is particularly true if use is made of spreadsheet solvers.

Let us illustrate the point by means of the game reserve planning problem. A spreadsheet formulation of the basic problem structure is displayed in Figure 1. The X_{sc} variables are stored in B4:D5, and the Y_{sj} variables in B6:D8. With reference to the constraints defined in Section 2:

- the first three constraints are represented by B13:D13=B15:D15;

- the next four constraints are represented by J4:K5≥J8:K9; and

- the final constraint is represented by H13≥J13.

[Readers unfamiliar with conventions adopted in the *What's Best* software, which was used for these examples, may need to note that the =>= notation in the figures represents a $\geq$ constraint which is tight in the current solution. The >= notation represents a $\geq$ constraint which is slack in the current solution. The software may also display not>= with obvious meaning, but this outcome is not illustrated in the examples.]

Figure 2 illustrates the set up for the Chebychev GP solution. Goals for the four objectives are specified in F17:F20 (and repeated for convenience in K17:K20), while the weights

	A	B	C	D	E	F	G	H	I	J	K
1		Numbers	(decision	variables):			Food Consumptions:			Food Available	
2			Species:					Species:			
3		A	B	C			A	B	C	Camp1	Camp2
4	Camp1	3.64	0.92	0.95		Stock1	1.00	3.00	8.00	14.00	25.00
5	Camp2	8.60	0.00	2.05		Stock2	2.00	1.00	4.00	12.00	30.00
6	Meat	1.20	0.28	0.00						=>=	=>=
7	Live Sales	3.69	0.00	0.16						=>=	>=
8	Hunting	0.00	0.00	0.44		Total consumed:	Stock1			14.00	25.00
9							Stock2			12.00	25.40
10											
11	Species Totals	12.24	0.92	3.00							
12	Growth Rate	0.40	0.30	0.20							
13	Production	4.90	0.28	0.60		Species C Total		3.00	=>=	3.00	
14		=	=	=							
15	Disposal	4.90	0.28	0.60							
16						Goals	Deviation		z_k+deviation		Goals
17	Revenues (z1)	11.09	0.833			13.00	1.91		13.00	=>=	13.00
18	Conservation (z2)	18.78	0.256			25.00	6.22		25.00	=>=	25.00
19	Local Food (z3)	28.17	0.050			60.00	31.83		60.00	=>=	60.00
20	Viewing (z4)	43.12	0.134			55.00	11.88		55.00	=>=	55.00
21											
22						Wt'd Deviations					
23	Epsilon weight	0.02					1.59	=<=			
24							1.59	=<=			
25	Total Objective:	1.72					1.59	=<=			
26							1.59	=<=			
27						D	1.59				

Figure 2: *Chebychev goal programming formulation for the game reserve planning LP.*

(w_i) are specified in C17:C20. The deviational variables δ_i are stored in G17:G20 and the weighted deviations $(w_i\delta_i)$ in G23:G26. The maximum deviation (Δ) is stored in G27. The additional constraints are implemented as follows:

- the deviation constraints (5) are represented by I17:I20≥K17:K20; and

- the maximum deviation constraints (7) are represented by G23:G26≤G27.

The objective function to be minimized is contained in B22.

Figure 3 illustrates the set up for the piecewise linear value function solution. The minimum values for each objective (z_i^0) are shown in F17:F20, and the variables z_{ik} in H17:K20. Values of $z_i^0 + \sum_{k=1}^{4} z_{ik}$ are stored in D17:D20, so that the constraints (11) are represented by B17:B20=D17:D20. Maximum segment lengths are defined in F23:F26, so that the upper bounds on the z_{ik} are represented by H17:K20≤$F$23:$F$26. The objective function parameters (β_{ik}) are stored in H30:K33, from which is calculated the overall value stored in B28 which is to be maximized.

It is clear that both approaches are thus simply formulated within a spreadsheet structure. We obtained the solutions by using the *What's Best* add-in from the LINDO corporation, but other solvers could equally well have been used.

The greater balance in the solutions (as compared to the use of naive linear aggregation) from both approaches is evident. These solutions are of course not numerically identical, as the preference information used in the examples for each method were separately specified for the purpose of illustration and are unlikely to be consistent. In both cases, good practice would be to conduct sensitivity analysis on this preference information, *i.e.* on the goal levels g_i in the case of goal programming, and on the weights W_i and the curvature of the value functions in the piecewise linear value function approach.

	A	B	C	D	E	F	G	H	I	J	K
1		Numbers	(decision variables):				Food Consumptions:			Food Available	
2			Species:					Species:			
3		A	B	C				A	B	C	Camp1	Camp2
4	Camp1	4.01	2.03	0.49		Stock1	1.00	3.00	8.00	14.00	25.00
5	Camp2	4.90	0.00	2.51		Stock2	2.00	1.00	4.00	12.00	30.00
6	Meat	0.44	0.00	0.00						=>=	=>=
7	LiveSales	3.12	0.00	0.00						=>=	>=
8	Hunting	0.00	0.61	0.60		Total consumed:	Stock1			14.00	25.00
9							Stock2			12.00	19.85
10											
11	Species Totals	8.91	2.03	3.00							
12	Growth Rate	0.40	0.30	0.20							
13	Production	3.56	0.61	0.60		Species C Total		3.00	=>=	3.00	
14		=	=	=							
15	Disposal	3.56	0.61	0.60							
16						min		Partial contributions			
17	Revenues (z1)	14.69	=	14.69		4.8		3.00	3.00	3.00	0.89
18	Conservation (z2)	15.60	=	15.60		0		7.80	7.80	0.00	0.00
19	Local Food (z3)	8.87	=	8.87		0		8.87	0.00	0.00	0.00
20	Viewing (z4)	55.33	=	55.33		33		11.17	11.17	0.00	0.00
21											
22						Segment length					
23						3.000		=<=	=<=	=<=	<=
24						7.800		=<=	=<=	<=	<=
25						30.000		<=	<=	<=	<=
26						11.167		=<=	=<=	<=	<=
27											
28	Overall Value:	47.171				Objective Function Construction					
29								10	7	4	2
30						10		3.3333	2.3333	1.3333	0.6667
31						8		1.0256	0.7179	0.4103	0.2051
32						6		0.2000	0.1400	0.0800	0.0400
33						6		0.5373	0.3761	0.2149	0.1075

Figure 3: *Piecewise linear value theory formulation for the game reserve planning LP.*

7 Conclusions

In this paper, we have argued firstly that good practice in any linear programming modelling should include identification of the inevitable multiplicity of objectives. Problem structuring for multiple objectives has been widely discussed in the MCDA literature (see, for example, Belton and Stewart [1], Chapter 3, and perhaps more comprehensively Rosenhead and Mingers [8]).

We have then demonstrated that a simple additive aggregation of the objectives is an extremely poor strategy. On the other hand there exist alternative strategies for dealing with multiobjective LP, based on goal programming and on piecewise linear value function approximations respectively. These strategies are easily implemented in standard LP structures without the need for specialized software. What is needed is for the analyst to commit more time and resources to assessing underlying preference structures of the decision maker, but this, surely, is the essence of OR.

References

[1] Belton V & Stewart TJ, 2002, *Multiple criteria decision analysis: An integrated approach*, Kluwer Academic Publishers, Boston (MA).

[2] Geldenhuys G, 1963, *Notes on some important techniques of operations research*, Technical report, National Research Institute for Mathematical Sciences, South African Council for Scientific and Industrial Research, Pretoria.

[3] HILLIER FS & LIEBERMAN GJ, 1995, *Introduction to operations research*, 6^{th} edition, McGraw-Hill, New York (NY).

[4] JORDI KC & PEDDIE D, 1988, *A wildlife management problem: A case study in multiple-objective linear programming*, Journal of the Operational Research Society, **39**, pp. 1011–1020.

[5] KAPLAN RS & NORTON DP, 1996, *The balanced scorecard*, Harvard Business School Press, Boston (MA).

[6] KEENEY RL & RAIFFA H, 1976, *Decisions with multiple objectives*, J. Wiley & Sons, New York (NY).

[7] RAGSDALE CT, 2004, *Spreadsheet modelling and decision analysis: A practical introduction to management science*, Thomson, Madison (OH).

[8] ROSENHEAD J & MINGERS J (Eds), 2001, *Rational analysis for a problematic world revisited*, 2^{nd} edition, John Wiley & Sons, Chichester.

[9] SIMON HA, 1976, *Administrative behavior*, 3^{rd} edition, The Free Press, New York (NY).

[10] STEWART TJ, 1995, *Simplified approaches for multi-criteria decision making under uncertainty*, Journal of Multi-Criteria Decision Analysis, **4**, pp. 246–258.

[11] STEWART TJ, 1996, *Robustness of additive value function methods in MCDM*, Journal of Multi-Criteria Decision Analysis, **5**, pp. 301–309.

[12] WIERZBICKI AP, 1999, *Reference point approaches*, Chapter 9 in GAL T, STEWART TJ & HANNE T (Eds), *Multicriteria decision making: Advances in MCDM models, algorithms, theory, and applications*, Kluwer Academic Publishers, Boston (MA).

[13] WINSTON WL, 2004, *Operations research: Applications and algorithms*, 4^{th} edition, Thomson – Brooks/Cole, Belmont (CA).