COWLES COMMISSION DISCUSSION PAPER: ECONOMICS NO. 2114

The Linear Team: An Example of Linear Programming Under Uncertainty[1]

Roy Radner

November 2, 1954

Introduction

In a previous paper, [1] the team problem was discussed in the case of a quadratic payoff function. Here we want to investigate the team that is generated by introducing uncertainty into a linear programming problem.[2] It will be assumed that the reader is familiar with the notion of a team as described in pages i and ii of the introduction to [1].

In the typical linear programming (from now on abbreviated to l.p.) problem, we are required to maximize a linear function of the finite dimensional vector a, subject to the constraint that a lie in some convex set C. Imagine now that, instead of a constant vector c, we have a random vector γ, and that the coordinates of a are determined by decision functions α_i, so that

$$a_i = \alpha_i(y_i)$$

1/ Research undertaken by the Cowles Commission for Research in Economics under contract Nonr-358(01), NR 047-006 with the Office of Naval Research

2/ Thanks to D. Bratton, J. Marschak and M. Beckmann for helpful discussions on the subject of this paper.

where y_i are given information variables. For a given information structure we want to choose $\alpha = (\alpha_1, \ldots, \alpha_n)$ so as to maximize the expected (gross) payoff

$$(1) \qquad U(\alpha) = E\,\gamma'\alpha(y)$$

subject to the constraint

$$(2) \qquad \alpha(y) \in C \quad \text{a.e.}$$ [1]

Using the usual definitions of addition and scalar multiplication of functions, the set of possible decision functions becomes a linear space (exactly what space will be made precise in the next section); (1) defines a linear function on this space, and (2) constrains η to lie in a convex set of functions. Thus the linear team problem can be regarded as an l.p. problem in the (generally infinite dimensional) space of functions. It should be added that the constraint set C can also be made random in certain ways, without destroying the l.p. character of the problem.

In section I the above remarks are made precise. In section II a special type of information structure is discussed. Section III shows how the usual analysis by means of Lagrangean multipliers leads to systems of random implicit prices. Some further remarks about the constraints are made in section IV.

1/ a.e. ≡ "almost everywhere", i.e. with probability one.

I. The Linear Team Problem as an L.P. Problem in Infinitely Many Dimensions

We now proceed to make precise the problem formulation indicated in the introduction.

The Information Structure:

Let X denote the basic probability measure space, with elements x, and let $Y_1,\dots,Y_n$ be n other probability spaces with probability measures induced by the n measurable functions $\eta_1,\dots,\eta_n$, respectively. Thus

$$y_i = \eta_i(x)$$

where the y_i (in Y_i) are called the information variables, and the η_i and Y_i taken together are called the information structure.

The Decision Functions:

For any given information structure, let A be the set of all $\alpha = (\alpha_1,\dots,\alpha_n)$ such that α_i is in L_2 of Y_i (i.e. $E\,[\alpha_i(y_i)]^2 < \infty$). The elements of A will be called decision functions; they are indeed functions which take on values in n-dimensional Euclidean space R^n. (The restriction that any decision function must have finite mean square is discussed below.)

The Payoff

Let G be the set of all $\gamma = (\gamma_1,\dots,\gamma_n)$ such that γ_i is in L_2 of X. The γ_i will be the random coefficients in the linear payoff function. Thus we will be maximizing

$$(1) \qquad U(\alpha) = E \sum_i \alpha_i(y_i)\, \gamma_i(x)$$

subject to certain constraints on α. Since $U(\alpha)$ is clearly a bounded linear functional on A, it can be represented as an inner product of α with some element of A; this element is in fact $\bar{\gamma} = (\bar{\gamma}_1, \ldots, \bar{\gamma}_n)$ where

$$\bar{\gamma}_i(y_i) = E\,(\gamma_i(x)|y_i)$$

so that (1) can be rewritten

$$(2) \qquad U(\alpha) = (\alpha, \bar{\gamma}) = E \sum_i \alpha_i(y_i)\, \bar{\gamma}(y_i)$$

<u>The Constraints</u>

For every x in X let $\mathcal{C}(x)$ be a closed convex set in R^n, and let $\mathcal{C}$ be the set of all α in A such that $\alpha(y)$ is in $C(x)$ a.e. $\mathcal{C}$ will be the constraint set for the decision functions. is obviously convex in A, but the following is not quite so obvious: ¶ <u>Lemma: $\mathcal{C}$ is closed in A.</u>

Let $\{\alpha^n\}$ be a sequence in $\mathcal{C}$ which converges in the mean to α (in A.) Consider the sequence $\{\alpha_1^n\}$ of the first "coordinates" of the α^n. By the Riesz-Fisher Theorem there is a subsequence $\{\beta_1^n\}$ of $\{\alpha_1^n\}$ which converges pointwise to α_1 a.e. Let $\{\beta^n\}$ be the corresponding subsequence of $\{\alpha^n\}$. Now consider the sequence $\{\beta_2^n\}$ of the second coordinates of the β^n, etc. In this manner one can construct a subsequence $\{\delta^n\}$ of $\{\alpha^n\}$ which converges pointwise (in the metric of R^n) to α a.e. But $C(x)$ is closed; hence if $\delta^n(x) \rightarrow \alpha(x)$ then $\alpha(x)$ is in $C(x)$; and thus $\alpha(x)$ is in $C(x)$ a.e. Q.E.D.

<u>The Problem:</u>

Given the information structure, maximize $u(\alpha)$ with respect to α, subject to α in $\mathcal{C}$.

The restriction that the decision functions α_i have finite mean square is a real one, as the following example will show. In this example, there is a maximizing a for every value of $\gamma(x)$, and (in a certain sense) a best decision function $\hat{\alpha}$, but neither $E\hat{\alpha}_1^2$ nor $u(\hat{\alpha})$ exist.

Example: Let C be the convex set in the plane bounded by the parabola $a_2 = a_1^2$, let X be the non-negative real numbers, let

$$\gamma_1(x) = \frac{2x}{x+1}$$

$$\gamma_2(x) = -\frac{1}{x+1}$$

and let $y_1 = y_2 = x$. It is easily verified that the line

$$\frac{2x}{x+1} a_1 - \frac{1}{x+1} a_2 = \frac{x^2}{x+1}$$

is tangent to the boundary of C at the point (x, x^2) so that for each x the best actions are

$$\hat{\alpha}_1(x) = x$$

$$\hat{\alpha}_2(x) = x^2$$

and the corresponding payoff is

$$u[\hat{\alpha}(x), x] = \frac{x^2}{x+1}$$

Now both γ_1 and γ_2 are bounded for $x \geq 0$, so that whatever the distribution of x, $E[\gamma_i(x)]^2 < \infty$. But

$$\frac{x^2}{x+1} = (x-1) + \frac{1}{x+1} \geq x - 1$$

so that for any distribution of x which has no finite moments, neither $E[\hat{\alpha}_1(x)]^2$ nor $E\, u[\alpha(x),x]$ exist.

In the above example the convex set C was not bounded, but in most situations that arise in practice there will be no harm done by making the

Assumption: C is bounded.

Under this assumption, any decision function α such that α_i is measurable on Y_i which satisfies

$$\alpha(y) \text{ in } C(x) \text{ a.e.}$$

is of course essentially bounded, and a fortiori is in A.

II. Information Variables With Independent Ranges

In this section we will study the consequences of the assumptions that the ranges of variation of the different information variables y_i are independent, and that the constraint set C is not random. and it will be seen that in this case the problem of finding best decision rules always reduces to a finite dimensional l.p. problem.

To define the "range" of a general probability measure in a useful way does not seem to be very easy. If, however, we assume that the given measure is absolutely continuous with respect to some other measure, then the range of the given measure can be defined as the set on which the probability density is greater than zero. With this definition in mind, we shall make the following

Assumption: For every i, the conditional range of y_i given all the other $y_j (j \neq i)$ is independent of the values of the other y_j.

For any α, and for each i, let $[\underline{a}_i, \bar{a}_i]$ be the smallest closed interval I such that

$$\text{Prob}\,[\alpha_i(y_i) \in I] = 1$$

It is easy to see that the effect of the above assumption is that the constraint $\alpha(y)$ in C a.e. (C is closed convex set) is equivalent to the requirement that the Cartesian product

$$R(\alpha) = \prod_i [\underline{a}_i, \bar{a}_i]$$

must be contained in C.

Given any particular rectangle R (i.e. a Cartesian product of intervals) we ask: what is the best α such that $R(\alpha) = R$? According to condition (N) on p. 6 of [1], if $\hat{\alpha}$ is the best decision rule for which $R(\alpha) = R$ then for every i and almost every y_i, $\hat{\alpha}_i(y_i)$ must maximize

$$\hat{\alpha}_i(y_i) \; E\,(\gamma_i | y_i) + \sum_{i \neq j} E\,(\hat{\alpha}_j(y_j)\,\gamma_j \mid y_i)$$

subject to $\underline{a}_i \leq \hat{\alpha}_i(y_i) \leq \bar{a}_i$. This gives:

$$\hat{\alpha}_i(y_i) = \begin{Bmatrix} \bar{a}_i \\ \underline{a}_i \end{Bmatrix} \text{ as } E\,(\gamma_i | y_i) \begin{Bmatrix} > \\ \leq \end{Bmatrix} 0.$$

Thus the best payoff corresponding to the rectangle R is

$$U(R) = \sum_i (\bar{a}_i \, \bar{d}_i + \underline{a}_i \, \underline{d}_i)$$

where
$$\bar{d}_i = E\,[E(\gamma_i | y_i) \mid E(\gamma_i | y_i) > 0] \;\; \text{Prob}\,[E\,(\gamma_i | y_i) > 0]$$

$$\underline{d}_i = E\,[E(\gamma_i | y_i) \mid E(\gamma_i | y_i) \leq 0] \;\; \text{Prob}\,[E\,(\gamma_i | y_i) \leq 0]$$

(Note that $\bar{d}_i \geq 0$ and $\underline{d}_i \leq 0$.)
Thus we have reduced our problem to one of maximizing $U(R)$, which is linear in the $\bar{a}_i$ and $\underline{a}_i$ subject to the condition that R be contained in C.

If C is defined by m linear constraints, (besides the constraints $a_i \geq 0$) we have reduced our problem to an ordinary linear programming problem in $2n$ variables (the $\bar{a}_i$ and a_i) with $(2^n)m$ linear constraints (besides the $2n$ positivity constraints). The number 2^n of course, gets very large, very quickly!

If in the original constraints

$$a\,T \leq b$$

the elements of T and b are all ≥ 0, then the above problem is greatly simplified. In that case it is not hard to show that all a_i must be zero. (This diagram will convince the reader immediately:)

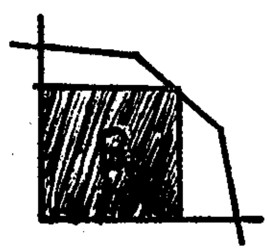

Thus the problem reduces to: "choose $\bar{a}_i$ so as to maximize

$$\Sigma\, \bar{d}_i\, \bar{a}_i$$

subject to $(\bar{a}_1, \ldots, \bar{a}_n)$ in C." This is an l.p. problem of the same "size" as the original, with the same constraints, and with coefficients which can be calculated from the distribution of the c_i's.

III. Random Lagrangean Multipliers and Implicit Prices

1. Introduction

Associated with an l.p. problem there is a non-negative saddle point problem involving the use of "Lagrangean multipliers"; furthermore these multipliers can often be interpreted as prices. In this section we will see how the multipliers for the team problem can be interpreted as random prices.

The application of the multiplier analysis to general linear spaces involved the extension of the Minkowski-Farkas Lemma to such spaces, and this seems to present certain mathematical difficulties. These difficulties will be avoided in this paper by assuming that the basic probability space X has only a finite number of elements. This implies that the space A of decision functions (see Sec. I) is finite dimensional.1/

We will assume that the convex set $\mathcal{C}$ in A is defined by the linear inequalities

$$(1) \qquad \left.\begin{array}{l} \alpha(y) \geq 0 \\ T(x)\,\alpha(y) \leq \bar{\beta}(x) \end{array}\right\} \text{ for all } x$$

where for each x, $T(x)$ is an $m \times n$ matrix, and $\bar{\beta}(x)$ is an m-dimensional vector.

Let B be the space of all m-dimensional vector valued functions of x. The equation

$$\beta(x) = T(x)\,\alpha(y)$$

defines a linear transformation τ from A to B. Condition (1) can thus be rewritten:

1/ For a discussion of the l.p. problem in general linear spaces see the recent discussion paper [3] by L. Hurwicz.

$$(2) \qquad \begin{cases} \alpha \geq 0 \\ \tau\alpha \leq \bar{\beta} \end{cases}$$

Since we will need to know the form of the adjoint of τ, we will settle this matter now in the following lemma.

Lemma. The adjoint τ^* of τ is given by:[1]

$$[\tau^* \beta]_j (y_j) = E \left\{ \sum_i t_{ij}(x) \beta_i(x) \mid y_j \right\}$$
$$= \left\{ [T'(x) \beta(x)]_j \mid y_j \right\}$$

Proof. For any x in A and β in B, τ^* must satisfy

$$(\alpha, \tau^*\beta) = (\tau\alpha, \beta)$$

$$E \sum_j \alpha_j(y_j) [\tau^*\beta]_j(y_j) = E \sum_i \sum_j t_{ij}(x) \alpha_j(y_j) \beta_i(x)$$

$$= E \sum_j \alpha_j(y_j) \sum_i t_{ij}(x) \beta_i(x)$$

Thus for any α_j

$$E \alpha_j(y_j) [\tau^*\beta]_j(y_j) = E \alpha_j(y_j) \sum_i t_{ij}(x) \beta_i(x)$$

$$= E \left(\alpha_j(y_j) E \left\{ \sum_i t_{ij}(x) \beta_i(x) \mid y_j \right\} \right)$$

Hence for every y_j

$$[\tau^*\beta]_j (y_j) = E \left\{ \sum_i t_{ij}(x) \beta_i(x) \mid y_j \right\} \qquad \text{Q.E.D.}$$

1/ If v is a vector, $[v]_j$ denotes the j'th coordinate of v. Thus if v is a vector valued function $[v]_j$ (x) denotes the j'th coordinate evaluated at x.

2. The Non-Negative Saddle Point Problem for the Linear Team.

Applying the well known saddle point theorem of linear programming to the l.p. problem just set up (see, for example [2], p. 487), we get immediately the following result:

The decision rule α^o maximizes

$$(\alpha, \bar{\gamma})$$

subject to the constraints (2) if and only if there is a $\lambda^o \geq 0$ in B such that the pair α^o, λ^o constitute a non-negative saddle point for the bilinear functional

(3) $$\phi(\alpha,\lambda) = (\alpha, \bar{\gamma}) + (\bar{\beta},\lambda) - (\tau\alpha, \lambda) = (\alpha,\bar{\gamma}) + (\bar{\beta},\lambda) - (\alpha, \tau^*\lambda)$$

i.e. for all $\alpha \geq 0$ and $\lambda \geq 0$

$$\phi(\alpha, \lambda^o) \leq \phi(\alpha^o, \lambda^o) \leq \phi(\alpha^o, \lambda)$$

Equivalently, α^o and λ^o satisfy the conditions:

(4a) $$\alpha^o \geq 0 \quad , \tau\alpha^o \leq \beta, \quad E(\bar{\gamma}(y) - (\tau^*\lambda)(y))'\,\alpha^o(y) = 0$$

(4b) $$\lambda^o \geq 0 \quad , \tau^*\lambda^o \geq \bar{\gamma}\ , \quad E\,(\bar{\beta}(x) - (\tau\alpha)(x)\,)'\,\lambda^o = 0$$

The transformation τ^* is given in the previous lemma.

The <u>dual problem</u> is: choose λ in B so as to minimize

$$E \sum_i \bar{\beta}_i(x)\, \lambda_i(x)$$

subject to the constraints

$$\begin{cases} \lambda_i(x) \geq 0 \\ E\left\{\sum_i t_{ij}(x)\,\lambda_i(x) \mid y_j\right\} \geq \bar{\gamma}_j(y_j) \end{cases}$$

3. Random Implicit Pricing

Just as in the case of certainty, the Lagrangean multipliers may be thought of as prices;[1] but being functions of x, they will be random prices. Thus suppose that for each x, $\sum \gamma_j(x)\, a_j$ represents the profit to a firm for operating the n activities at levels a_j, and that $t_{ij}(x)$ represents the added amount of input i $(i = 1,\ldots,m)$ needed to increase the level of the j'th activity by one unit. Suppose further that there is a manager for each activity, a custodian for each input, and a "helmsman", or central planning board.[2] Let $\lambda_i(x)$ be interpreted as the "shadow price" of the i'th input, for each value of x. From (3) we see that each manager must choose a rule α_j for setting the level of his activity which will maximize his shadow profit

$$E\left\{\alpha_j(y_j)\, E\,[\gamma_j(x) - \sum_i t_{ij}(x)\,\lambda_j(x) \mid y_j]\right\}$$

That is to say, he will choose an α_j with the property that:

$$\alpha_j(y_j)\begin{Bmatrix} > \\ = \end{Bmatrix} 0, \quad \text{according as} \quad E\,[\gamma_j(x) - \sum_i t_{ij}(x)\,\lambda_i(x) \mid y_j]\begin{Bmatrix} > \\ \leq \end{Bmatrix} 0$$

Koopmans (ibid.) has given the following rules of behavior for the helmsman, custodians and managers, to be applied in the ordinary (non-random) problem:

1/ See Koopmans, [4], pp. 93-95, for the certainty case.

2/ Each manager may actually be in charge of several activities; to bring this into our scheme, it is only necessary to require that all the corresponding information variables be identical.

"I. For the helmsman: Choose a vector p_{fin} of positive prices on all final commodities, and inform the custodian of each such commodity of its price.

II. For all custodians: Buy and sell your commodity from and to managers at one price only, which you announce to all managers. Buy all that is offered at that price. Sell all that is demanded up to the limit of availability.

III. For all custodians of final commodities: Announce to managers the price set on your commodity by the helmsman.

IV. For all custodians of intermediate commodities: Announce a tentative price on your commodity. If demand by managers falls short of supply by managers, lower your price. If demand exceeds supply, raise it.

V. For all custodians of primary commodities: Regard the available inflow from nature as a part of the supply of your commodity. Then follow the rule on custodians of intermediate commodities, with the following exception: Do no announce a price lower than zero but accept a demand below supply at a zero price if necessary.

VI. For all managers: Do not engage in activities that have negative profitability. Maintain activities of zero profitability at a constant level. Expand activities of positive profitability by increasing orders for the necessary inputs with, and offers of the outputs in question to, the custodians of those commodities.

"The dynamic aspects of these rules have on purpose been left vague. It is not specified by how much managers of profitable activities should increase their orders, or by how much custodians of commodities in short or excess supply should change the price. Neither is it indicated how during a temporary disequilibrium a commodity in short supply is apportioned to managers. These questions would be highly relevant if our purpose were to design an allocation model which automatically seeks and finds an efficient point from some initial nonoptimal situation. However, our present purpose is only to demonstrate that an efficient point, once achieved, is maintained if all players follow the rules stated."

Before these rules can be applied to the case of uncertainty, the following additions and amendments must be made.

I'. For the helmsman: Choose a vector valued function λ indicating that the set of prices $\lambda(x)$ will be effective if the state of nature which is actually realized is x. Inform the custodians and managers of this price function λ.

II'. For the jth custodian: Sell your input at the price $\lambda_j(x)$ corresponding to the state of nature actually realized.

VI. For all managers: The term "profitability" must be interpreted as "expected profitability given the information y_j (for manager j).

The dynamics under uncertainty are even vaguer than under certainty. Suppose the process is thought of as consisting of repeated trials, with a new state of nature x, being realized at each trial. Should prices and activity levels be changed after each trial, or should time be allowed in which to estimate the average or expected results from each policy? Such problems as these would have to be studied (in addition to the ones already inherent in the certainty case) before this analysis could be used as the basis for an "optimum-seeking model".

IV. The Role of the Constraints in an Uncertain Program.

Because of our requirement that the given constraints be satisfied with certainty, even though the information on which the activity levels are based is random, the convex set $\mathcal{C}$ of decision rules may be quite small in many instances. This corresponds to our intuitive feeling that if different decisions are based on <u>different</u> (random) information, then it may be very difficult to make sure that the constraints are always satisfied, without unduly restricting the range over which the decisions can vary. We will illustrate this by a two variable example.

Let the linear form be:

$$U = E[\gamma_1 \alpha_1(y_1) + \gamma_2 \alpha_2(y_2)]$$

and the constraints be:

$$\alpha_1(y_1) + \alpha_2(y_2) \leq 1, \quad \alpha_i(y_i) \geq 0$$

Suppose that the conditional ranges of y_1 and y_2 are independent, as in Sec. II. The result at the end of that section applies here, and so we get that

$$\alpha_i(y_i) = \begin{Bmatrix} \bar{a}_i \\ 0 \end{Bmatrix} \text{ as } E(\gamma_i | y_i) \begin{Bmatrix} > \\ \leq \end{Bmatrix} 0$$

where $\bar{a}_1$ and $\bar{a}_2$ must be chosen to maximize

$$\bar{d}_1 \bar{a}_1 + \bar{d}_2 \bar{a}_2$$

subject to

$$\bar{a}_1 + \bar{a}_2 \leq 1, \quad \bar{a}_i \geq 0$$

where $\bar{d}_i = E[E(\gamma_i | y_i) \mid E(\gamma_i | y_i) > 0] \text{ Prob } [E(\gamma_i | y_i) > 0]$

It is easy to see that for the best choice, $\bar{a}_j = 1$ for that j for which $\bar{d}_j$ is the largest, and the other $\bar{a}_i = 0$. Thus the level of only one of the activities actually depends upon the information received.

Since the constraints in an l.p. problem usually arise as approximations in a situation in which the loss incurred by violating the constraints rises sharply with the magnitude of the violation, the disturbing situation exemplified above might be corrected by dropping the constraints and modifying the payoff function appropriately. This of course would change the problem into an ordinary maximum problem, without the l.p. "flavor". In some cases, however, the loss incurred by violating the constraints might be realistically represented as the cost of some linear program needed to remedy the violation. Such a situation would be covered by our formulation, by dividing the activities in two groups: those which are based on incompleted information, but are not subject to constraints; and those which are based on complete information (the remedial activities) and are subject to constraints.

References

[1] R. Radner, "The Quadratic Team. I.", CCDP: Statistics No. 386, Aug. 2, 1954. (hectographed)

[2] H. Kuhn and A. Tucker, "Nonlinear Programming", in Second Berkeley Symposium on Mathematical Statistics and Probability (Neyman, ed.), U. of California Press, Berkeley, 1951 (p. 481-492).

[3] L. Hurwicz, "Programming in General Spaces", CCDP: Economics No. 2109, July 29, 1954 (hectographed)

[4] T. Koopmans, "Analysis of Production as an Efficient Combination of Activities", Ch. III of Activity Analysis of Production and Allocation, Wiley, N.Y., 1951.