A LAGRANGIAN FINITE GENERATION TECHNIQUE
FOR SOLVING LINEAR-QUADRATIC PROBLEMS
IN STOCHASTIC PROGRAMMING

R.T. Rockafellar
R. J.-B. Wets

April 1984
WP-84-25

A LAGRANGIAN FINITE GENERATION TECHNIQUE FOR SOLVING LINEAR-QUADRATIC PROBLEMS IN STOCHASTIC PROGRAMMING

R. T. Rockafellar* and R.J.-B. Wets*

ABSTRACT

A new method is proposed for solving two-stage problems in linear and quadratic stochastic programming. Such problems are dualized, and the dual, although itself of high dimension, is approximated by a sequence of quadratic programming subproblems whose dimensionality can be kept low. These subproblems correspond to maximizing the dual objective over the convex hull of finitely many dual feasible solutions. An optimizing sequence is produced for the primal problem that converges at a linear rate in the strongly quadratic case. An outer algorithm of augmented Lagrangian type can be used to introduce strongly quadratic terms, if desired.

1. Introduction

In the recourse model in stochastic programming, a vector x must be chosen optimally with respect to present costs and constraints as well as certain expected costs and induced constraints that are associated with corrective actions available in the future. Such actions may be taken in response to the observation of the values of various random variables about which there is only statistical information at the time x is selected. The actions involve costs and constraints that depend on these observed values and on x. The theory of this kind of stochastic programming and the numerical methods that have been proposed for it has been surveyed recently by Wets [10].

We aim here at developing a new solution procedure for the case where the first and second stage problems in the recourse model fit the mold of linear or quadratic (convex) programming. We assume for simplicity that the random variables are discretely distributed with only finitely many values. This restriction is not fully necessary in theory, but it reflects the realities of computation and a natural division among the questions that arise. Every continuous distribution must in practice be replaced by a finite discrete one, whether empirically, or through sampling, mathematical approximation, or in connection with the numerical calculation of integrals expressing expectations. The effects of such discretization raise important questions of convergence and statistical confidence in the solutions that are obtained, but such matters are best left to separate study.

We assume therefore that the probability space is a finite set Ω: the probability associated with an element $\omega \in \Omega$ is p_ω, and the expectation of a quantity u_ω that depends on ω is

$$E_\omega u_\omega := \sum_{\omega \in \Omega} p_\omega u_\omega.$$

The fundamental problem we want to address is

(P) $\qquad$ minimize $c \cdot x + \frac{1}{2} x \cdot C x + E_\omega \psi_\omega(x)$ over all $x \in X \subset R^n$,

where X is a nonempty convex polyhedron, c is a vector in R^n, C is a symmetric matrix in $R^{n \times n}$ that is positive semidefinite, and $\psi_\omega(x)$ is the minimum cost in a certain recourse subproblem that depends on ω and x. We view this recourse subproblem as one of linear or quadratic programming, but instead of handling it directly we work with its dual. More will be said about this later (Proposition 1 in §2), but what counts in the end is the following: we suppose a representation

(1.1) $\qquad \psi_\omega(x) = \max_{z_\omega \in Z_\omega} \{ z_\omega \cdot [h_\omega - T_\omega x] - \frac{1}{2} z_\omega \cdot H_\omega z_\omega \}$

is on hand, where Z_ω is a nonempty convex polyhedron in R^m, T_ω is a matrix in $R^{m \times n}$, h_ω is a vector in R^m, and H_ω is a symmetric matrix in $R^{m \times m}$ that is positive semidefinite. Note from the subscript ω that all the elements in this representation are in principle allowed to be random, although a particular application might not involve quite so much randomness.

Two basic conditions are imposed on the given data. We assume X and C are such that for every $v \in R^n$ the set

$$(1.2) \qquad \xi(v) := \operatorname*{argmin}_{x \in X} \{ v \cdot x + \tfrac{1}{2} x \cdot C x \}$$

is nonempty and bounded. We also assume Z_ω, h_ω, T_ω, and H_ω, are such that for every $x \in X$ the set

$$(1.3) \qquad \zeta_\omega(x) := \operatorname*{argmax}_{z_\omega \in Z_\omega} \{ z_\omega \cdot [h_\omega - T_\omega x] - \tfrac{1}{2} z_\omega \cdot H_\omega z_\omega \}$$

is nonempty and bounded. Certainly the first condition holds if X is bounded or C is positive definite, and the second holds if Z_ω is bounded or H_ω is positive definite.

The first condition is quite innocuous, since in practice X can always be taken to be bounded. It implies that the function

$$(1.4) \qquad \varphi(v) = \inf_{x \in X} \{ v \cdot x + \tfrac{1}{2} x \cdot C x \},$$

which will have a role in duality, is finite everywhere.

The second condition is more subtle, since it involves dual elements that might not be given directly but derived instead from a primal statement of the recourse subproblem that depends on x and ω. It ensures in particular that for every $x \in X$ and $w \in \Omega$, the optimal value $\psi_\omega(x)$ in this subproblem is finite, and an optimal recourse exists. This means that our stochastic programming problem (P) is one of *relatively complete recourse* [8]: there are no induced constraints on x that arise from the need to keep open the possibility of recourse at a later time.

Of course, if our problem were not one of relatively complete recourse, we could make it so by identifying the induced constraints and shrinking the set X until they were all satisfied. The smaller X would still be a convex polyhedron, although its description might be tedious. In this sense our second condition forces no real restriction on the problem either, except in requiring that the induced constraints, if any, be identified thoroughly in advance.

The term $E_\omega \psi_\omega(x)$ in (P) can also be interpreted as an expected *penalty*, incidentally. In terms of

$$(1.5) \qquad \theta_\omega(u) = \max_{z_\omega \in Z_\omega} \{ z_\omega \cdot u - \tfrac{1}{2} z_\omega \cdot H_\omega z_\omega \}$$

we have

$$(1.6) \qquad \psi_\omega(x) = \theta_\omega(h_\omega - T_\omega x).$$

If $0 \in Z_\omega$, then

$$(1.7) \qquad \theta_\omega(u) \geq 0 \text{ for all } u, \quad \theta_\omega(0) = 0.$$

We can view $\theta_\omega(h_\omega - T_\omega x)$ as a penalty attached to certain degrees or directions of deviation of $T_\omega x$ from the vector h_ω. Many suitable penalty functions can be expressed as in (1.5), and this provides further motivation for taking $\psi_\omega(x)$ to be of the form (1.1). Note that the case where

$\theta_\omega(h_\omega - T_\omega x)$ is a sum of separate terms, one for each real component of the deviation vector $h_\omega - T_\omega x$, can be identified with the case where each Z_ω is a product of intervals and H_ω is diagonal.

The solution procedure that we shall present depends on a Lagrangian representation of problem (P) which leads to the dual problem

(D)   maximize $\varphi(c - E_\omega T^*_\omega z_\omega) + E_\omega \{z_\omega \cdot h_\omega - \frac{1}{2} z_\omega \cdot H_\omega z_\omega\}$

subject to $z_\omega \in Z_\omega$ for all $\omega \in \Omega$.

Here φ is the function in (1.4), for which another representation will later be given (Proposition 2 in §2). The asterisk $*$ signals the transpose of a matrix. The maximization in (D) takes place over the convex polyhedron

(1.8) $Z = \Pi_{\omega \in \Omega} Z_\omega \subset (R^m)^\Omega;$

we think of z_ω as the component in Z_ω of a point $z \in Z$. The vector space $(R^m)^\Omega$ here, which is a produce of copies of R^m, one for each $\omega \in \Omega$, is likely to be of very high dimension, since the number of points in Ω may be very large. Despite this formidable dimensionality it is by way of (D), at least in concept, that we propose to solve (P). Properties of expectation, decomposition and quadratic structure, will make this plausible. The relationship between (P) and (D) is explored in §2 along with other issues of quadratic programming duality that are important in our formulation.

We approach problem (D) by a "finite generation" technique in which the feasible region Z is approximated from within by polytopes of comparatively low dimension, a polytope being a subset generated as the convex hull of finitely many points. This technique is presented in §3. It resembles the classical finite-element or Galerkin approach to the unconstrained maximization of a functional defined over an infinite-dimensional space, where one maximizes over finite-dimensional subspaces that grow in size as the approximation is refined. An important difference, however, is that in our case the new element or elements that are introduced at each stage in modifying the polytope over which we maximize are not obtained from some predetermined scheme, as classically, but identified in an "adaptive" manner. Furthermore, the total number of elements used in generating the polytope does not have to keep increasing; the sequence of polytopes does not have to be nested. We prove in §4 that when the matrix C is positive definite these elements can readily be consolidated without threat to ultimate convergence, although the rate of progress may be better if a substantial set of generating elements is maintained. In this way the dimension of the subproblem to be solved in every iteration can be kept as low as seems desirable.

The subproblem of maximizing over a polytope can be represented as a standard type of quadratic programming problem and solved exactly by available codes. It yields as a byproduct an approximate solution vector for (P) along with bounds that provide a test of near optimality. The sequence of such approximate solutions converges to an optimal solution to (P). If not only C but also the matrices H_ω are positive definite, the rate of convergence is linear, in fact with guaranteed progress of a

certain sort in every iteration, not just for the tail of the sequence.

In producing a new element to be used in the polytopal subrepresentation of Z, we have a particular x on hand and must carry out the maximization in (1.1) for every $\omega \in \Omega$. In other words, we must solve a large number of closely related linear or quadratic programming problems in R^m. This could be a difficult task in general, but techniques such as have already been developed in connection with other approaches to stochastic programming problems of a more special nature (see Wets [10]) do offer hope. Furthermore, there are cases of definite interest where the maximization in (1.1) is trivial, for instance where Z_ω is a product of intervals and H_ω is diagonal. Such a case has been described in [9].

Not all of the problems we wish to solve have C and H_ω positive definite, but this does not prevent the application of our method and the achievement of a linear rate of convergence. Augmented Lagrangian techniques [5] can be effective in approximating any problem (P) by a sequence of similar problems that do exhibit positive definiteness. We explain this in §5 after having established in §4 the results that show the advantages of the strongly quadratic case.

2. Lagrangian Representation and Duality.

As the *Lagrangian* associated with problem (P) under the representation (1.1) of the recourse costs, we shall mean the function

$$(2.1) \quad L(x,z) = c \cdot x + \tfrac{1}{2} x \cdot Cx + E_\omega \{ z_\omega \cdot [h_\omega - T_\omega x] - \tfrac{1}{2} z_\omega \cdot H_\omega z_\omega \} \quad \text{for} \quad x \in X, z \in Z,$$

where Z is the convex polyhedron in (1.5). Clearly $L(x,z)$ is convex in x and concave in z, since C and H_ω are positive semidefinite. General duality theory [6] associates with L, X, and Z, the primal problem

$$(2.2) \qquad \text{minimize } F \text{ over } X, \text{ where } F(x) := \max_{z \in Z} L(x,z),$$

and the dual problem

$$(2.3) \qquad \text{maximize } G \text{ over } Z, \text{ where } G(z) := \min_{x \in X} L(x,z).$$

The functions F and G are convex and concave, respectively. Our assumptions in §1 allow us to write "max" and "min" in their definitions rather than "sup" and "inf".

These problems turn out to be the ones already introduced. In terms of the notation in (1.2) and (1.3), we have

$$(2.4) \qquad \underset{z \in Z}{\text{argmax}}\, L(x,z) = \{ z \mid z_\omega \in \zeta_\omega(x) \text{ for all } \omega \in \Omega \},$$

$$(2.5) \qquad \underset{x \in X}{\text{argmin}}\, L(x,z) = \xi(c - E_\omega T_\omega^* z_\omega).$$

Moreover for $x \in X$ and $z \in Z$ we have

$$(2.6) \qquad F(x) = c \cdot x + \tfrac{1}{2} x \cdot Cx + E_\omega \psi_\omega(x),$$

$$(2.7) \qquad G(z) = \varphi(c - E_\omega T_\omega^* z_\omega) + E_\omega \{ z_\omega \cdot h_\omega - \tfrac{1}{2} z_\omega \cdot H_\omega z_\omega \}.$$

Thus the primal and dual problems (2.2) and (2.3) can be identified with (P) and (D), respectively.

In order to continue with our analysis of these problems, we need to step back briefly for a look at some basic facts about duality in quadratic programming, not only as they might apply to (P) and (D), but also to various subproblems in our schemes. A quadratic programming problem is usually defined as a problem in which a quadratic convex function is minimized (or a quadratic concave function maximized) subject to a system of linear constraints, or in other words, over a convex polyhedron. As is well known, such a problem has an optimal solution whenever its optimal value is finite (see Frank and Wolfe [3, Appendix (i)]); the Kuhn-Tucker conditions are both necessary and sufficient for optimality. For the purpose at hand, it is essential to adopt a more general point of view in which a problem is considered to fall in the category of quadratic programming as long as it can be *represented* in this traditional form, possibly through the introduction of auxiliary variables.

Consider an arbitrary Lagrangian of the form

$$(2.8) \qquad l(u,v) = p \cdot u + q \cdot v + \tfrac{1}{2} u \cdot Pu - \tfrac{1}{2} v \cdot Qv - v \cdot Ru \quad \text{for} \quad u \in U, v \in V,$$

where U and V are nonempty convex polyhedra, and P and Q are symmetric, positive semidefinite matrices. Let

(2.9) $$f(u) = \sup_{v \in V}\{v \cdot [q - Ru] - \tfrac{1}{2}v \cdot Qv\},$$

(2.10) $$U_0 = \{u \mid f(u) \text{ finite}\} = \{u \mid \sup \text{ in } (2.9) \text{ attained}\},$$

(2.11) $$g(v) = \inf_{u \in U}\{u \cdot [p - R^*v] - \tfrac{1}{2}u \cdot Pu\},$$

(2.12) $$V_0 = \{v \mid g(v) \text{ finite}\} = \{v \mid \inf \text{ in } (2.11) \text{ attained}\}.$$

The primal and dual problems associated with l, U, and V by general duality theory can then be written as:

(P$_0$) minimize $p \cdot u + \tfrac{1}{2}u \cdot Pu + f(u)$ over $u \in U \cap U_0$,

(D$_0$) maximize $q \cdot v - \tfrac{1}{2}v \cdot Qv + g(v)$ over $v \in V \cap V_0$.

The following duality theorem for (P$_0$) and (D$_0$) extends the standard results in quadratic programming that were achieved by Dorn [2] and Cottle [1]. Those authors concentrated in effect on the case where U and V are orthants.

THEOREM 1. *Problems* (P$_0$) *and* (D$_0$) *are representable as quadratic programming in the traditional sense. If* (P$_0$) *and* (D$_0$) *both have feasible solutions, or if either* (P$_0$) *or* (D$_0$) *has finite optimal value, then both have optimal solutions, and*

$$\min(P_0) = \max(D_0).$$

This occurs if and only if the Lagrangian l has a saddle point $(\overline{u},\overline{v})$ relative to $U \times V$, in which case the saddle value $l(\overline{u},\overline{v})$ coincides with the common optimal value in (P$_0$) *and* (D$_0$), *and the saddle points are the pairs $(\overline{u},\overline{v})$ such that $\overline{u}$ is an optimal solution to* (P$_0$) *and $\overline{v}$ is an optimal solution to* (D$_0$).

PROOF. General duality theory [6] assures us that $\inf(P_0) \geq \sup(D_0)$ and in particular that both (P$_0$) and (D$_0$) have finite optimal value if both have feasible solutions. It also informs us that $(\overline{u},\overline{v})$ is a saddle point of l on $U \times V$ if and only if $\overline{u}$ is an optimal solution to (P$_0$), $\overline{v}$ is an optimal solution to (D$_0$), and $\min(P_0) = \max(D_0)$, this common optimal value then being equal to $l(\overline{u},\overline{v})$. We know further that a quadratic programming problem in the traditional sense has an optimal solution if it has finite optimal value [3, Appendix (i)]. The Kuhn-Tucker conditions are both necessary and sufficient for optimality in such a problem, because the constraint system is linear. The proof of the theorem can be reduced therefore to demonstrating that (P$_0$) and (D$_0$) are representable as quadratic programming in the traditional sense and in such a manner that the Kuhn-Tucker conditions for either problem correspond to the saddle point condition for l on $U \times V$.

The sets U and V are associated with systems of linear constraints that can be expressed in various ways, but to be specific we can suppose that

(2.13) $U = \{u \in R^n \mid Au \geq a\} \neq \phi$ and $V = \{v \in R^m \mid B_v^* \leq b\} \neq \phi$,

where A is $m' \times n$ and B is $m \times n'$. Let $u' \in R^{n'}$ and $v' \in R^{m'}$ be Lagrange multiplier vectors paired with the conditions $B^* v \leq b$ and $Au \geq a$, respectively.

Formula (2.9) gives $f(u)$ as the optimal value in a classical quadratic programming problem in v. The optimal solutions to this problem are vectors that satisfy the usual Kuhn-Tucker conditions, or in other words, correspond to saddle points of the Lagrangian

(2.14)
$$v \cdot [q - Ru] - \tfrac{1}{2} v \cdot Qv + u' \cdot [b - B^* v]$$

$$= b \cdot u' + v \cdot [q - Ru - Bu'] - \tfrac{1}{2} v \cdot Qv$$

relative to $u' \in R_+^{n'}$ and $v \in R^m$. In particular, then, we have

(2.15) $f(u) = \inf_{u' \in R_+^{n'}} \sup_{v \in R^m} \{b \times u' + v \cdot [q - Ru - Bu'] - \tfrac{1}{2} v \cdot Qv\}$.

The inner supremum here is attained whenever finite, and it is attained at a point $v = u''$. Thus it equals ∞ unless there exists a vector $u'' \in R^m$ such that $[q - Ru - Bu'] - Qu'' = 0$, in which case it equals $b \cdot u' + \tfrac{1}{2} u'' \cdot Qu''$, a value that actually depends only on u and u'. We may conclude that

(2.16) $U_0 = \{u \in R^n \mid \exists u' \in R_+^{n'}, \text{ with } Ru + Bu' + Qu'' = q\}$,

(2.17) $f(u) = $ minimum of $b \cdot u' + \tfrac{1}{2} u'' \cdot Qu''$

subject to $u' \in R_+^{n'}, u'' \in R^m, \quad Ru + Ru' + Qu'' = q$.

We can therefore represent (P_0) as

$(\tilde{P}_0)$ minimize $p \cdot u + \tfrac{1}{2} u \cdot Pu + b \cdot u' + \tfrac{1}{2} u'' \cdot Qu''$

subject to $Au \geq a, u' \geq 0, Ru + Bu' + Qu'' = q$,

where the value of $u'' \cdot Qu''$ does not depend on the particular choice of the vector u'' satisfying $Ru + Bu' + Qu'' = q$ but only on u and u'. This is a quadratic programming problem in the usual sense, but in which u'' is a sort of vector of dummy variables that can be eliminated, if desired. In any case it follows that (P_0) has an optimal solution if its optimal value is finite, inasmuch as this property holds for $(\tilde{P}_0)$.

The optimal solutions $(\overline{u}, \overline{u}', \overline{u}'')$ to $(\tilde{P}_0)$ are characterized by the Kuhn-Tucker conditions that involve multiplier vectors $\overline{v}$ for the constraint $Ru + Bu' + Qu'' = q$ and $\overline{v}'$ for the constraint $Au \geq a$. These conditions take the form:

$$A\overline{u} \ge a, \quad \overline{v} \ge 0, \quad \overline{v} \cdot [A\overline{u} - a] = 0,$$

$$\overline{u} \ge 0, \quad B^*\overline{v} \le b, \quad \overline{u} \cdot [B^*\overline{v} - b] = 0,$$

$$R^*\overline{v} + A^*\overline{v}' - P\overline{u} = p, \quad R\overline{u} + B\overline{u}' + Q\overline{u}'' = q, \quad Q\overline{u}'' = Q\overline{v}.$$

Because of the final condition we can write the next-to-last condition instead as $R\overline{u} + B\overline{u}' + Q\overline{v} = q$. Note that there is no restriction then on $\overline{u}''$, except that $Q\overline{u}'' = Q\overline{v}$; we always take $\overline{u}'' = \overline{v}$ in particular. This is in keeping with our observation that (P_0) is really just a problem in u and u'. We see in fact that the pairs $(\overline{u}, \overline{u}')$ which are optimal for (P_0) are the ones which, for some pair $(\overline{v}, \overline{v}')$, satisfy the conditions

(2.18)
$$A\overline{u} \ge 0, \quad \overline{v} \ge 0, \quad \overline{v} \cdot [A\overline{u} - a] = 0,$$

$$\overline{u} \ge 0, \quad B^*\overline{v} \le b, \quad \overline{u} \cdot [B^*\overline{v} - b] = 0,$$

$$R\overline{u} + B\overline{u}' + Q\overline{v} = q, \quad R^*\overline{v} + A^*\overline{v}' - P\overline{u} = p.$$

Problem (D_0) can be understood in the same way. From the formula (2.11) for $g(v)$ we deduce that

(2.19)
$$V_0 = \{v \in R^m \mid \exists v' \in R_+^{m'}, \ v'' \in R^N, \ \text{with} \ R^*v + B^*v' - Pv'' = p\},$$

(2.20)
$$g(v) = \text{maximum of} \ a \cdot v' - \tfrac{1}{2}v'' \cdot Pv'' \ \text{subject to}$$

$$v' \in R_+^{m'}, \ v'' \in R^n, \ R^*v + A^*v' - Pv'' = p.$$

These formulas yield for (D_0) the representation

$(\tilde{D}_0)$ maximize $q \cdot v - \tfrac{1}{2}v \cdot Qv + a \cdot v' - \tfrac{1}{2}v'' \cdot Pv''$ subject to

$$B^*v \le b, \ v' \ge 0, \ R^*v + A^*v' - Pv'' = p,$$

where the value of $v'' \cdot Pv''$ does not depend on the particular v'' satisfying $R^*v + A^*v' - Pv'' = p$ but only on v and v'. This is really a problem in v and v', and the Kuhn-Tucker conditions characterize $\overline{v}$ and $\overline{v}'$, as optimal if and only if there exist $\overline{u}$ and $\overline{u}'$ such that (2.18) holds, the same conditions as before. Since $(\tilde{D}_0)$ is a quadratic programming problem in the usual sense, it has an optimal solution whenever its optimal value is finite, and (D_0) therefore has this property too.

Our argument demonstrates that if either (P_0) or (D_0) has finite optimal value, then both problems have optimal solutions. The optimal solutions in both cases are characterized by the existence of auxiliary vectors such that (2.18) holds. But (2.18) can also be seen as the Kuhn-Tucker conditions for $(\overline{u}, \overline{v})$ to be a saddle point of the Lagrangian (2.8), when U and V are given by (2.13). Thus for $\overline{u}$ and $\overline{v}$ to be optimal solutions to (P_0) and (D_0) respectively, it is necessary and sufficient that $(\overline{u}, \overline{v})$ be a saddle point in (2.18). Following on the remarks at the beginning of the proof, this establishes the theorem. []

COROLLARY. *Any standard quadratic programming method can in principle be used to solve problems of the form* (P_0) *or* (D_0), *in fact both*

simultaneously, thereby determining a saddle point of the corresponding Lagrangian l on $U \times V$, if such a saddle point exists.

PROOF. The representations in the proof of the theorem show more specifically that if an algorithm is applied to (P_0), the optimal solution vectors $\overline{u}, \overline{u}'$ and multiplier vectors $\overline{v}$ and $\overline{v}'$ which it produces yield optimal solutions $\overline{u}$ to (P_0) and $\overline{v}$ to (D_0), and $(\overline{u}, \overline{v})$ is a saddle point in (2.8). The same holds if an algorithm is applied to (D_0), except that then $\overline{v}$ and $\overline{v}'$ are the optimal solution vectors, whereas $\overline{u}$ and $\overline{u}'$ are the multiplier vectors. []

THEOREM 2. *The stochastic programming problems (P) and (D) are representable as quadratic programming problems in the traditional sense, although with potentially very high dimensionality. Both problems have optimal solutions, and*

$$\min(P) = \max(D).$$

A pair $(\overline{x}, \overline{z})$ is a saddle point of the Lagrangian L relative to $X \times Z$ if and only if $\overline{x}$ is an optimal solution to (P) and $\overline{z}$ is an optimal solution to (D). The set of such pairs $(\overline{x}, \overline{z})$ is bounded.

PROOF. We need only observe that the triple L, X, Z, can be construed as a special case of the triple l, U, V, in Theorem 1. A term like $E_\omega z_\omega \cdot H_\omega z_\omega$ can be expressed as $z \cdot Qz$ for certain matrix Q, and so forth. Our assumption that the extremal sets $\xi(v)$ in (1.2) and $\zeta_\omega(x)$ in (1.3) are nonempty for all $v \in R^n$, $x \in X$ and $\omega \in \Omega$, guarantees that every $x \in X$ is feasible for (P), and every $z \in Z$ is feasible for (D). Therefore we are in the case of Theorem 1 where both problems have feasible solutions.

As for the boundedness of the set of saddle points $(\overline{x}, \overline{z})$, consider a particular pair of optimal solutions $\overline{x}^*$ and $\overline{z}^*$ to (P) and (D). Observe that for every optimal solution $\overline{x}$ to (P), $(\overline{x}, \overline{z}^*)$ is a saddle point and therefore satisfies

$$\overline{x} \in \operatorname*{argmin}_{x \in X} L(x, \overline{z}^*) = \xi(c - E_\omega T_\omega^* \overline{z}_\omega^*)$$

(cf. (2.5)). But the set on the right is bounded (one of our basic assumptions in §1). Likewise for every optimal solution $\overline{z}$ to (D), $(\overline{x}^*, \overline{z})$ is a saddle point and therefore satisfies

$$\overline{z} \in \operatorname*{argmax}_{z \in Z} L(\overline{x}^*, z), \quad \text{so} \quad \overline{z}_\omega \in \zeta_\omega(\overline{x}^*) \quad \text{for all} \quad \omega \in \Omega.$$

(cf. (2.4)). The sets $\zeta_\omega(\overline{x}^*)$ are all bounded (again by one of our basic assumptions in §1), so $\overline{z}$ belongs to a certain bounded set. The pairs $(\overline{x}, \overline{z})$ thus all belong to a product of bounded sets dependent only on $\overline{x}^*$ and $\overline{z}^*$. []

The following pair of results will help to clarify the quadratic programming nature of problems (P) and (D).

PROPOSITION 1. *For the function ψ_ω given by (1.1), if the polytope Z_ω has a representation*

(2.21) $$Z_\omega = \{z_\omega \in R^m \mid B_\omega^* z_\omega \leq b_\omega\}$$

for some vector $b_\omega \in R^s$ and matrix $B_\omega \in R^{m \times s}$ (with s independent of ω), then ψ_ω has an alternative expression of the form

(2.22) $$\psi_\omega(x) = minimum \ of \ \ d_\omega \cdot y_\omega + \frac{1}{2} y_\omega \cdot D_\omega y_\omega$$

$$subject \ to \ \ y_\omega \in Y_\omega, \ \ \overline{T}_\omega x + W_\omega y_\omega = \overline{h}_\omega,$$

for certain vectors $d_\omega \in R^s$, $\overline{h}_\omega \in R^q$, and matrices $\overline{T}_\omega \in R^{q \times n}$, $W_\omega \in R^{q \times s}$, and $D_\omega \in R^{s \times s}$ with D_ω symmetric and positive semidefinite, and where

(2.23) $$Y_\omega = \{y_\omega \in R^s \mid A_\omega y_\omega \geq a_\omega\}$$

for some $a_\omega \in R^p$ and $A_\omega \in R^{p \times s}$.

Conversely, any function ψ_ω having a representation (2.22) as just described (with $\psi_\omega(x)$ finite for all $x \in X$) also has a representation (1.1) with Z_ω of the form (2.21).

PROOF. Starting with the representation (1.1) and Z_ω of the form (2.21), view the maximization problem in (1.1) as the dual problem associated with the Lagrangian

$$l_{x,\omega}(u_\omega, z_\omega) = u_\omega \cdot [b_\omega - B_\omega^* z_\omega] + z_\omega \cdot [h_\omega - T_\omega x] - \frac{1}{2} z_\omega \cdot H_\omega z_\omega$$

$$for \ u_\omega \in R_+^s \ and \ z_\omega \in R^m.$$

The corresponding primal problem, whose optimal value is also equal to $\psi_\omega(x)$ by Theorem 1 (as long as $x \in X$, so that $\psi_\omega(x)$ is finite by assumption) is

$$minimize \ b_\omega \cdot u_\omega + f_\omega(u_\omega) \ over \ u_\omega \in R_+^s, \ where$$

$$f_\omega(u_\omega) = \sup_{z_\omega \in R^m} \{z_\omega \cdot [h_\omega - T_\omega x - B_\omega u_\omega] - \frac{1}{2} z_\omega \cdot H_\omega z_\omega\}.$$

Using the trick in the proof of Theorem 1, we can reformulate the latter as

$$minimize \ b_\omega \cdot u_\omega + \frac{1}{2} u_\omega'' \cdot H_\omega u_\omega'' \ subject \ to$$

$$u_\omega \in R_+^s, \ u_\omega'' \in R^m, \ B_\omega u_\omega + H_\omega u_\omega'' = h_\omega - T_\omega x.$$

We can then pass to form (2.22) in terms of $y_\omega = (u_\omega, u_\omega'')$ (or by setting $y_\omega = u_\omega$ after algebraic elimination of u_ω'', if the rank of H_ω is the same for all $\omega \in \Omega$).

Starting with the representation (2.22) and Y_ω of the form (2.23), on the other hand, we can view $\psi_\omega(x)$ as the optimal value for the primal problem associated with the Lagrangian

$$l_{x,w}(y_\omega,v_\omega) = d_\omega \bullet y_\omega + \tfrac{1}{2} y_\omega \bullet D_\omega y_\omega + v_\omega \bullet [\overline{h}_\omega - \overline{T}_\omega x - W_\omega y_\omega]$$

for $y_\omega \in Y_\omega$ and $v_\omega \in R^q$.

Then $\psi_\omega(x)$ (when finite) is also the optimal value in the corresponding dual problem

maximize $v_\omega \bullet [\overline{h}_\omega - \overline{T}_\omega x] + g_\omega(v_\omega)$ over $v_\omega \in R^q$, where

$$g_\omega(v_\omega) = \inf_{y_\omega \in Y_\omega} \{ y_\omega \bullet [d_\omega - W_\omega^* v_\omega] + \tfrac{1}{2} y_\omega \bullet D_\omega y_\omega \}.$$

As we saw in the proof of Theorem 1, this problem can also be written as

maximize $v_\omega \bullet [\overline{h}_\omega - \overline{T}_\omega x] + v_\omega' \bullet a_\omega - \tfrac{1}{2} v_\omega'' \bullet D_\omega v_\omega''$ subject to

$$v_\omega \in R^q, \quad v_\omega' \in R^p_+, \quad W_\omega^* v_\omega + A_\omega^* v_\omega' + D_\omega v_\omega'' = d_\omega.$$

With $z_\omega = (v_\omega, v_\omega', v_\omega'')$, this can be brought into the form (1.1) with Z_ω as in (2.21). (Alternatively one could take $z_\omega = (v_\omega, v_\omega')$ and eliminate v_ω'' algebraically, provided that the rank of D_ω is independent of ω. If the rank of W_ω is also independent of ω, one could even eliminate v_ω from the problem and just take $z_\omega = v_\omega'$ to get a representation (1.1) in fewer variables.) ▯

PROPOSITION 2. *The function φ in (1.4) also has a representation*

(2.24) $\varphi(v) = maximum\ of\ q \bullet u - \tfrac{1}{2} u \bullet Qu\ over\ all$

$u \in U\ satisfying\ Bu = v$

for some choice of vectors b and q and matrices B and Q with Q symmetric and positive semidefinite, where U is a convex polyhedron.

PROOF.   Recall that $\varphi(v)$ is finite for all v by assumption. Express X as $\{x \in R^n \mid Ax \geq a\}$ for some $a \in R^p$ and $A \in R^{p \times n}$, and consider the Lagrangian

$$l_v(x,u') = v \bullet x + \tfrac{1}{2} x \bullet Cx + u' \bullet [a - Ax] \text{ for } x \in R^n \text{ and } u' \in R^p_+.$$

The primal problem associated with this Lagrangian is the minimization problem in (1.4), whereas the dual problem, which also has $\varphi(v)$ as its optimal value, is

maximize $a \bullet u' + g(u')$ over $u' \in R^p_+$, where

$$g(u') = \inf_{x \in R^n} \{ x \bullet [v - A^* u'] + \tfrac{1}{2} x \bullet Cx \}.$$

The reformulation trick in Theorem 1 translates this into

maximize $a \bullet u' - \tfrac{1}{2} u'' \bullet Cu''$ subject to

$$u' \in R^p_+, \quad u'' \in R^n, \quad A^* u' - Cu'' = v.$$

We can then get a representation (2.24) in terms of $u = (u',u'')$. $\square$

Propositions 1 and 2 make possible a more complete description of the quadratic programming representation of problems (P) and (D) indicated in Theorem 2. When $\psi_\omega(x)$ is expressed in terms of a recourse subproblem in y_ω as in Proposition 1, we can identify (P) with the problem

$$(2.25) \qquad \text{minimize } c \cdot x + \tfrac{1}{2} x \cdot Cx + E_\omega\{d_\omega \cdot y_\omega + \tfrac{1}{2} y_\omega \cdot D_\omega y_\omega\}$$

$$\text{subject to } x \in X, \; y_\omega \in Y_\omega, \; \overline{T}_\omega x + W_\omega \, y_\omega = \overline{h_\omega} \text{ for all } \omega \in \Omega.$$

Similarly, when φ is expressed as in Proposition 2 we can pose (D) as

$$(2.26) \qquad \text{maximize } q \cdot u - \tfrac{1}{2} u \cdot Qu + E_\omega\{z_\omega \cdot h_\omega - \tfrac{1}{2} z_\omega \cdot H_\omega z_\omega\}$$

$$\text{subject to } u \in U, \; z_\omega \in Z_\omega, \text{ and } Bu + E_\omega\{T_\omega^* z_\omega\} = c.$$

In the latter, our assumption that $\varphi(v)$ is finite for all $v \in R^n$ implies that no matter what the choice of vectors $z_\omega \in Z_\omega$, there does exist a $u \in U$ such that the constraint $Bu + E_\omega\{T_\omega^* z_\omega\} = c$ is satisfied.

3. Finite Generation Method

Our aim is to solve problem (P) by way of (D) according to the following scheme. We replace (D) by a sequence of subproblems

(D^ν) maximize $G(z)$ over all $z \in Z^\nu \subset Z$

for $\nu = 1, 2, \ldots$, where G is the dual objective function in (2.3) and (2.7), and Z^ν is a polytope of relatively low dimension generated as the convex hull of finitely many points in Z. Obviously (D^ν) is the dual of

(P^ν) minimize $F^\nu(x)$ over all $x \in X$,

where

$$(3.1) \qquad F^\nu(x) = \max_{z \in Z^\nu} L(x, z)$$

$$= c \cdot x + \tfrac{1}{2} x \cdot Cx + \max_{z \in Z^\nu} E_\omega \{ z_\omega \cdot [h_\omega - T_\omega x] - \tfrac{1}{2} z_\omega \cdot H_\omega z_\omega \}.$$

Indeed, (P^ν) and (D^ν) are the primal and dual problems that correspond to L on $X \times Z^\nu$ rather than $X \times Z$. In calculating a solution $\bar{z}^\nu$ to (D^ν) we obtain also a solution $\bar{x}^\nu$ to (P^ν) that can be viewed as an approximately optimal solution to (P). From $\bar{z}^\nu$ and $\bar{x}^\nu$ we gain information that determines the polytope $Z^{\nu+1}$ to be used in the next iteration. The new polytope $Z^{\nu+1}$ is not necessarily "larger" than Z^ν.

Problems (P^ν) and (D^ν) belong to the realm of generalized quadratic programming as demarcated in §2. Clearly

$$(3.2) \qquad F(x) \geq F^\nu(x) \text{ for all } x,$$

where F is the primal objective function in (2.2) and (2.6), so (P^ν) can be regarded as a "lower envelope approximation" to (P). The feasible sets in (P^ν) and (D^ν) are the same as in (P) and (D), namely X and Z. From Theorem 1, therefore, we know that optimal solutions $\bar{x}^\nu$ and $\bar{z}^\nu$ to (P^ν) and (D^ν) exist and satisfy

$$(3.3) \qquad F^\nu(\bar{x}^\nu) = G(\bar{z}^\nu),$$

$$(3.4) \qquad \bar{x}^\nu \in \operatorname*{argmin}_{\bar{x} \in X} F^\nu(x) \subset \operatorname*{argmin}_{\bar{x} \in X} L(x, \bar{z}^\nu),$$

$$(3.5) \qquad \bar{z}^\nu \in \operatorname*{argmax}_{z \in Z^\nu} G(z) \subset \operatorname*{argmax}_{z \in Z^\nu} L(\bar{x}^\nu, z).$$

Having determined such a pair $(\bar{x}^\nu, \bar{z}^\nu)$, which is a saddle point of L relative to $X \times Z^\nu$, we can test whether it is actually a saddle point of L relative to $X \times Z$. This amounts to checking the maximum of $L(\bar{x}^\nu, z)$ over all $z \in Z$ to see if it occurs at $z = \bar{z}^\nu$. If yes, $\bar{x}^\nu$ and $\bar{z}^\nu$ are optimal solutions to (P) and (D), and we are done. If no, we obtain from the test an element

$$(3.6) \qquad z^\nu \in \operatorname*{argmax}_{z \in Z} L(\bar{x}^\nu, z)$$

and have

$$(3.7) \qquad L(\bar{x}^\nu, \bar{z}^\nu) < L(\bar{x}^\nu, z^\nu) = F(\bar{x}^\nu).$$

The crucial feature that makes the test possible is the decomposition in (2.4): maximizing $L(\bar{x}^\nu, z)$ in $z \in Z$ reduces to solving a separate quadratic

programming problem (perhaps trivial) in $z_\omega \in Z_\omega$ for each $\omega \in \Omega$. Anyway, with such a z^ν we have

(3.8) $F(x) \geq L(x, z^\nu)$ for all x, with equality when $x = \bar{x}^\nu$.

We can use this in conjunction with (3.3) in constructing a new lower envelope approximation $F^{\nu+1}$ for F, which in primal terms is what is involved in constructing a new set $Z^{\nu+1}$ to replace Z^ν. More will be said about this later.

Of course the optimality test also furnishes a criterion for termination with suboptimal solutions, if desired. Since $\bar{x}^\nu$ and $\bar{z}^\nu$ are feasible solutions to (P) and (D) with

(3.9) $F(\bar{x}^\nu) \geq \min(P) = \max(D) \geq G(\bar{z}^\nu)$,

we know that for $\varepsilon_\nu = F(\bar{x}^\nu) - G(\bar{z}^\nu)$, $\bar{x}^\nu$ and $\bar{z}^\nu$ are ε_ν-optimal:

(3.10) $|F(\bar{x}^\nu) - \min(P)| \leq \varepsilon_\nu$ and $|G(\bar{x}^\nu) - \max(D)| \leq \varepsilon_\nu$

Our basic procedure can be summarized now as follows.

ALGORITHM.

Step 0 (Initialization). Choose the optimality test parameter $\bar{\varepsilon} \geq 0$ and the initial convex polytope $Z^1 \subset Z$. Set $\nu = 1$.

Step 1 (Approximate Solution). Determine a saddle point $(\bar{x}^\nu, \bar{z}^\nu)$ of L relative to $X \times Z^\nu$ and the value $\bar{\alpha}_\nu = L(\bar{x}^\nu, \bar{z}^\nu)$.

Step 2 (Direction Search). For each $\omega \in \Omega$, determine an optimal solution z_ω^ν to the problem

(3.11) $\text{maximize } z_\omega \cdot [h_\omega - T_\omega \bar{x}^\nu] - \frac{1}{2} z_\omega \cdot H_\omega z_\omega \text{ over } z_\omega \in Z_\omega$

and the optimal value α_ω^ν. Let z^ν be the element of Z having component z_ω^ν in Z_ω, and let

(3.12) $\alpha_\nu = c \cdot \bar{x}^\nu + \frac{1}{2} \bar{x}^\nu \cdot C \bar{x}^\nu + E_\omega \alpha_\omega^\nu = L(\bar{x}^\nu, z^\nu)$.

Step 3 (Optimality Test). Let $\varepsilon_\nu = \alpha_\nu - \bar{\alpha}_\nu$. Then $\bar{x}^\nu$ is an ε_ν-optimal solution to (P), $\bar{z}^\nu$ is an ε_ν-optimal solution to (D), and

(3.13) $\alpha_\nu \geq \min(P) = \max(D) \geq \bar{\alpha}_\nu$.

If $\varepsilon_\nu \leq \bar{\varepsilon}$, terminate.

Step 4 (Polytope Modification). Choose a new convex polytope $Z^{\nu+1}$ that contains both $\bar{z}^\nu$ and z^ν, although not necessarily all of Z^ν. Replace ν by $\nu+1$ and return to Step 1.

We proceed to comment on these steps in more detail, one by one.

The most important observation concerns the quadratic programming nature of the subproblem solved in Step 1. Suppose that Z^ν is generated from certain elements $\tilde{z}_k^\nu \in Z$:

(3.14) $Z^\nu = \text{co}\{\tilde{z}_k^\nu \mid k = 1, \ldots, m_\nu\}$

$$= \{\textstyle\sum_{k=1}^{m_\nu}\lambda_k \tilde{z}_k^\nu \mid \lambda_k \geq 0, \sum_{k=1}^{m_\nu}\lambda_k = 1\}.$$

Finding a saddle point $(\bar{x}^\nu,\bar{z}^\nu)$ of $L(x,z)$ relative $x\in X$ and $z\in Z^\nu$ is equivalent to finding a saddle point $(\bar{x}^\nu,\bar{\lambda}^\nu)$ of

$$(3.15) \qquad L^\nu(x,\lambda) = L(x,\textstyle\sum_{k=1}^{m_\nu}\lambda_k\tilde{z}_k^\nu)$$

relative to $x\in X$ and $\lambda\in\Lambda^\nu$, where Λ^ν is the unit simplex in R^{m_ν},

$$(3.16) \qquad \Lambda^\nu := \{\lambda=(\lambda_1,\dots,\lambda_{m_\nu}) \mid \lambda_k\geq 0, \textstyle\sum_{k=1}^{m_\nu}\lambda_k = 1\},$$

and then setting

$$(3.17) \qquad \bar{z}^\nu = \textstyle\sum_{k=1}^{m_\nu}\lambda_k\tilde{z}_k^\nu.$$

But from the definition (2.1) of $L(x,z)$ we have

$$(3.18) \qquad L^\nu(x,\lambda) = c\bullet x +\tfrac{1}{2}x\bullet Cx + \textstyle\sum_{k=1}^{m_\nu}E_\omega\{\tilde{z}_{k\omega}^\nu\bullet[h_\omega - T_\omega x]\}$$

$$-\tfrac{1}{2}\textstyle\sum_{j=1}^{m_\nu}\sum_{k=1}^{m_\nu}\lambda_j\lambda_k E_\omega\{\tilde{z}_{k\omega}^\nu\bullet H_\omega\tilde{z}_{k\omega}^\nu\}$$

$$= c\bullet x +\tfrac{1}{2}x\bullet Cx + \lambda\bullet[\tilde{h}^\nu - \tilde{T}^\nu x] -\tfrac{1}{2}\lambda\bullet\tilde{H}^\nu\lambda,$$

where

$$(3.19) \qquad \tilde{h}^\nu\in R^{m_\nu} \text{ with components } \tilde{h}_k^\nu = E_\omega\{\tilde{z}_{k\omega}^\nu\bullet h_\omega\},$$

$$(3.20) \qquad \tilde{H}^\nu\in R^{m_\nu\times m_\nu} \text{ with entries } \tilde{H}_{jk}^\nu = E_\omega\{\tilde{z}_{j\omega}^\nu\bullet H_\omega\tilde{z}_{k\omega}^\nu\},$$

$$(3.21) \qquad \tilde{T}^\nu\in R^{m_\nu\times n} \text{ with entries } \tilde{T}_{ki}^\nu = E_\omega\{\tilde{z}_{k\omega}^\nu\bullet T_\omega^i\},$$

T_ω^i being the ith column of the matrix $T_\omega\in R^{m\times m}$.

Problem (D^ν) thus reduces to a *deterministic quadratic programming problem in which the coefficients are certain expectations*, namely

$$(\tilde{D}^\nu) \qquad \text{maximize } \varphi(c - \tilde{T}^{\nu*}\lambda) -\tfrac{1}{2}\lambda\bullet\tilde{H}^\nu\lambda \text{ over all } \lambda\in\Lambda^\nu.$$

Here φ is the function in (1.4), which has alternative representations such as in Proposition 2 that can be used to place $(\tilde{D}^\nu)$ in a more traditional quadratic programming format. Regardless of such reformulation, the dimensionality of this quadratic programming problem will be relatively low as long as m_ν, the number of elements $\tilde{z}_k^\nu$ used in generating Z^ν, is kept modest.

The translation of (D^ν) into $(\tilde{D}^\nu)$ also sheds light on the lower envelope function F^ν in the approximate primal subproblem (P^ν):

$$(3.22) \qquad F^\nu(x) = \max_{\lambda\in\Lambda^\nu} L(x,\lambda) = c\bullet x +\tfrac{1}{2}x\bullet Cx + \Psi^\nu(x),$$

where

$$(3.23) \qquad \Psi^\nu(x) := \max_{\lambda\in\Lambda^\nu}\{\lambda\bullet[\tilde{h}^\nu - \tilde{H}^\nu x] -\tfrac{1}{2}\lambda\bullet\tilde{H}^\nu\lambda\}$$

$$= \max_{z\in Z^\nu} E_\omega\{z_\omega\bullet[h_\omega - T_\omega x] -\tfrac{1}{2}z_\omega\bullet H_\omega z_\omega\} = E_\omega\psi_\omega(x).$$

Clearly $\Psi^\nu(x)$ is a lower envelope approximation to the recourse cost function

$$(3.24) \qquad \Psi(x) := \max_{z \in Z} E_\omega \{ z_\omega \bullet [h_\omega - T_\omega x] - \tfrac{1}{2} z_\omega \bullet H_\omega z_\omega \} = E_\omega \psi_\omega(x).$$

Especially worth noting in (3.23) is the case where there are no quadratic terms $z_\omega \bullet H_\omega z_\omega$, i.e. where $H_\omega = 0$ for all $\omega \in \Omega$ and consequently $\tilde{H}^\nu = 0$. Then

$$(3.25) \qquad \Psi^\nu(x) = \max_{k=1,\ldots,m_\nu} \{ \tilde{h}_k^\nu - \tilde{T}_k^\nu \bullet x \},$$

where $\tilde{T}_k^\nu$ is the vector in R^n given by the k^{th} row of the matrix $\tilde{T}^\nu$ in (3.21):

$$(3.26) \qquad \tilde{T}_k^\nu = E_\omega \{ T_\omega^* \; \tilde{z}_{k\omega}^\nu \}.$$

In this case Ψ^ν is a *polyhedral* convex envelope representation of Ψ, the pointwise maximum of a collection of affine functions

$$l_k(x) = \tilde{h}_k^\nu - \tilde{T}_k^\nu \bullet x \text{ for } k=1,\ldots,m_\nu.$$

Our technique then resembles a cutting-plane method, at least as far as the function Ψ is concerned.

Indeed, if not only $H_\omega = 0$ but $C = 0$, so that there are no quadratic cost terms at all and (P) is a purely linear stochastic programming problem, we can regard F^ν as a polyhedral convex representation of F. Then the subproblems (P^ν) and (D^ν) can be solved by linear rather than quadratic programming algorithms. Furthermore the function $L(x,z^\nu)$ determined in (3.8) is then affine in x. If we were to take $Z^{\nu+1} = \text{co}\{Z^\nu, z^\nu\}$, we would get

$$F^{\nu+1}(x) = \max\{F^\nu(x), L(x,z^\nu)\},$$

and this would truly be a cutting-plane method applied to problem (P).

It must be remembered, though, that in such a cutting-plane approach it would generally be necessary to retain more and more affine functions in the polyhedral approximation to F. The dimension of the linear programming subproblem to be solved in each iteration would become progressively larger. In contrast, by taking advantage of the quadratic structure even to the extent of introducing it when it is not already at hand (as proposed in §5), one can avoid the escalation of dimensionality and at the same time get convergence results of a superior character.

Note that with a nonvanishing quadratic term $\lambda \bullet \tilde{H}^\nu \lambda$ in (3.23) (the matrix $\tilde{H}^\nu$ being positive semidefinite, of course) the lower envelope approximation Ψ^ν to Ψ will generally *not* be polyhedral but have "rounded corners". As a matter of fact, if $\tilde{H}^\nu$ is nonsingular, then Ψ^ν is a smooth convex function with Lipschitz continuous derivatives.

In Step 2 of the algorithm, we need to solve a potentially large number of quadratic programming problems (3.11) in the vectors z_ω. This could be a trouble spot. If the problems are complicated and require full application of some quadratic programming routine, the secret to

success would have to lie in taking advantage of the similarities between neighboring problems. Techniques of parametric programming and "bunching" might be useful. Not to be overlooked, however, are the situations in which each problem (3.11) decomposes further into something simpler. Especially important is the case where

$$(3.27) \qquad Z_\omega = Z_{\omega 1} \times Z_{\omega 2} \times \cdots \times Z_{\omega r}$$

and H_ω does not involve crossterms between the sets in this product:

$$(3.28) \qquad H_\omega = \text{diag}\,[H_{\omega 1}, H_{\omega 2}, ..., H_{\omega r}].$$

Then (3.11) reduces to a separate problem over each of the sets $Z_{\omega 1}, ..., Z_{\omega r}$. If these sets are actually intervals (bounded or unbounded), the separate problems are one-dimensional, and their solutions can be given in closed form. Such is indeed what happens when the costs $\psi_\omega(x)$ in (P) are penalties $\theta_\omega(h_\omega - T_\omega x)$ as in (1.5), (1.6), (1.7), and $\theta_\omega(h_\omega - T_\omega x)$ is a sum of separate terms, one for each real component of the vector $h_\omega - T_\omega x$.

The product case (3.27) also raises further possibilities for handling the subproblems in Step 1, by the way. We can write

$$(3.29) \qquad Z = Z_1 \times ... \times Z_r \text{ with } Z_j = \Pi_{\omega \in \Omega} Z_{\omega j}$$

and work with polytopes of the form

$$(3.30) \qquad Z^\nu = Z_1^\nu \times ... \times Z_r^\nu \text{ with } Z_j^\nu \subset Z_j,$$

for instance. This could be advantageous in holding the dimensionality down. If each Z_j^ν is generated as the convex hull of a finite subset of Z_j consisting of n_ν elements, we can get away with describing the points of Z^ν by rn_ν parameters λ_{jk}. On the other hand, if Z^ν is regarded as the convex hull of the product of these finite subsets of $Z_1, ..., Z_r$, we would need $(n_\nu)^r$ parameters.

The procedure invoked in Step 4 of the algorithm could be influenced by such considerations too. For this reason it has been left open to various possibilities. Two possibilities that immediately come to mind are:

$$(3.31) \qquad Z^{\nu+1} = \text{co}\,\{\bar{z}^\nu, z^\nu\} \quad \text{(generalized Frank-Wolfe rule)}$$

and

$$(3.32) \qquad Z^{\nu+1} = \text{co}\,\{Z^\nu, z^\nu\} \quad \text{(generalized cutting-plane rule)}.$$

The first of these is adequate for convergence if the matrix C is positive definite, as we shall see in §4. It is certainly the simplest but might suffer from too much information being thrown away between one iteration of Step 1 and the next. It gets its name from the interpretation in terms of problem (D) that will underly the proof of Theorem 5.

The second formula goes to the opposite extreme. It achieves better and better representations of the primal objective F, in the sense that

$$(3.33) \qquad F(x) \geq F^{\nu+1}(x) \geq \max\{F^\nu(x), L(x, z^\nu)\} \text{ for all } x,$$

$$\text{with } F(\bar{x}^\nu) = F^{\nu+1}(\bar{x}^\nu, z^\nu),$$

but this is at the expense of keeping all information and continually enlarging the size of the quadratic programming subproblem. A good compromise possibility is

$$(3.34) \qquad Z^{\nu+1} = \text{co}\{Z^1, \bar{z}^\nu, z^\nu\},$$

where Z^1 is the fixed initial polytope.

This brings us to the choice of Z^1 in Step 0, which in determining the first approximate solutions $\bar{x}^1$ and $\bar{z}^1$ could have a big effect on the progress of the computations. We can, of course, start with $Z^1 = \{\hat{z}\}$, where $\hat{z}$ is an element of Z that may be regarded as an estimate for an optimal solution to (D). For example, if an initial guess $\hat{x}$ is available for an optimal solution to (P), one might take $\hat{z}$ to be a vector constructed by calculating an element $\hat{z}_\omega \in \zeta_\omega(\hat{x})$ for each ω. This approach makes sense especially in situations where $\zeta_\omega(\hat{x})$ is a singleton for each $\omega \in \Omega$, so that $\hat{z}$ is uniquely determined by the estimate $\hat{x}$.

Another approach to the initial Z^1 requires no guesses or prior information about solutions. A fixed number of elements $a_{k\omega}(k=1,...,p)$ is chosen from each Z_ω, such as the set of extreme points of Z_ω augmented by some selected internal points. These yield p elements a_k of Z, where a_k has component $a_{k\omega}$ in Z_ω. The convex hull of these a_k's can be taken as Z^1.

Particularly interesting here is the case where Z_ω is a polytope independent of ω:

$$(3.35) \qquad Z_\omega = \text{co}\{a_1,...,a_p\} \subset R^m \text{ for all } \omega \in \Omega.$$

Then in taking Z^1 to be the convex hull of the corresponding "constant vectors" $a_k \in Z$ for these points $a_k \in R^m$, we get a very special form for the subproblem in Step 1 (cf. formulas (3.19), (3.20) and (3.21) with $\tilde{z}_{k\omega}^1 = a_k$ for all $\omega \in \Omega$). This subproblem $(\tilde{D}^1)$ is equivalent to the problem (D_1) obtained from (D) by restricting attention to the "constant vectors" z in Z (whose component in Z_ω is the same for all ω) and replacing h_ω, H_ω and T_ω accordingly by their expectations

$$(3.36) \qquad \tilde{h} = E_\omega h_\omega, \quad \tilde{H} = E_\omega H_\omega, \quad \tilde{T} = E_\omega T_\omega.$$

This idea could be refined further: we could partition Ω into subsets $\Omega_1,...,\Omega_l$, and restrict attention in (D) to vectors z whose component z_ω was constant in ω relative to each of these subsets. Correspondingly in (3.36) we would have conditional expectations. The resulting problem could again be identified with the (D_1) associated with the choice of a particular $Z^1 \subset Z$, namely the convex hull of the finitely many vectors z of the type just mentioned whose components all belong to the set $\{a_1, \ldots, a_p\}$. This Z^1 could be represented economically as a product set, and so forth.

In summary, there are many possibilities for choosing the initial polytope Z^1 and modifying it iteratively in Step 4. They can be tailored to the structure of the problem. Various product representations of Z and Z^ν could be helpful in particular. Versions of rules (3.31), (3.32), and (3.34), which maintain the product form can be developed.

See the end of §4 for other comments on forming $Z^{\nu+1}$ from Z^ν.

4. Convergence Results.

Properties of the sequences produced by the finite generation algorithm in §3 will now be derived. For this purpose we ignore the optimality test in Step 3 of the algorithm, since our interest is centered on what happens when the procedure is iterated indefinitely. Unless otherwise indicated, our assumptions are merely the basic ones in §1. The initial polytope Z^1 is arbitrary, and $Z^{\nu+1}$ is not subjected to any requirement stricter than the one in Step 4, namely that $Z^{\nu+1} \supset \{\bar{z}^\nu, z^\nu\}$. We use the supplementary notation

$$(4.1) \qquad \bar{\alpha} = \min(P) = \max(D),$$

$$(4.2) \qquad \bar{\varepsilon}_\nu = \bar{\alpha} - \bar{\alpha}_\nu,$$

$$(4.3) \qquad \bar{w}^\nu = \nabla_x L(\bar{x}^\nu, \bar{z}^\nu) = c + C\bar{x}^\nu - E_\omega T_\omega^* \bar{z}_\omega^\nu,$$

$$(4.4) \qquad \|x\|_C = [x \cdot Cx]^{1/2}.$$

Of course $\|\cdot\|_C$ is a norm on R^n if C is positive definite. If C is only positive semidefinite, then $\|x\|_C$ vanishes on the subspace $\{x \in R^n \mid Cx = 0\}$ but is positive elsewhere.

THEOREM 3. *The sequences $\{\bar{x}^\nu\}$, $\{\bar{z}^\nu\}$, and $\{z^\nu\}$ are bounded and satisfy*

$$(4.5) \qquad F(\bar{x}^\nu) = \alpha_\nu \geq \bar{\alpha} \geq \cdots \geq \bar{\alpha}_{\nu+1} \geq \bar{\alpha}_\nu = G(\bar{z}^\nu).$$

Furthermore one has the estimate

$$(4.6) \qquad \tfrac{1}{2}\|\bar{x} - \bar{x}^\nu\|_C^2 \leq \bar{\varepsilon}_\nu - \bar{w}^\nu \cdot (\bar{x} - \bar{x}^\nu) \leq \bar{\varepsilon}_\nu \leq \varepsilon_\nu$$

for every optimal solution $\bar{x}$ to (P), where

$$(4.7) \qquad \bar{w}^\nu \cdot (x - \bar{x}^\nu) \geq 0 \text{ for every } x \in X.$$

If $\varepsilon_\nu \to 0$, then every cluster point of $\{\bar{x}^\nu\}$ is an optimal solution to (P), and every cluster point of $\{\bar{z}^\nu\}$ is an optimal solution to (D).

PROOF. We have $\bar{\alpha}_\nu = L(\bar{x}^\nu, \bar{z}^\nu)$ and $\alpha_\nu = L(\bar{x}^\nu, z^\nu)$ by definition, so $F(\bar{x}^\nu) = \alpha_\nu$ by (3.7). Then $\alpha_\nu \geq \bar{\alpha} \geq \bar{\alpha}_\nu$ by (3.9). By the same token, $G(\bar{z}^{\nu+1}) = \bar{\alpha}_{\nu+1}$ and $\bar{\alpha} \geq \bar{\alpha}_{\nu+1}$. But also

$$G(\bar{z}^{\nu+1}) = \max_{z \in Z^{\nu+1}} G(z) \geq G(\bar{z}^\nu)$$

because $\bar{z}^\nu \in Z^{\nu+1}$. All the relations in (4.5) are therefore correct.

Next we verify that the sequence $\{\bar{z}^\nu\}$ is bounded. Recall that G is a continuous concave function on Z, since G is given by (2.7), where φ is the concave function defined by (1.4); our basic assumption about the sets $\xi(v)$ being bounded implies φ is finite elsewhere. (As is well known, a concave function is continuous at a point if it is finite on a neighborhood of the point [7, Theorem 10.1].) We know from (4.5) that the sequence $\{G(\bar{z}^\nu)\}$ is nondecreasing so the boundedness of $\{\bar{z}^\nu\}$ can be established by showing that the set $\{z \in Z \mid G(z) \geq G(\bar{z}^1)\}$ is bounded. Consider the closed concave function

$$g(z) = \begin{cases} G(z) & \text{if } z \in Z, \\ \infty & \text{if } z \notin Z. \end{cases}$$

We wish to show that a certain level is set $\{z \mid g(z) \geq \overline{\alpha}_1\}$ is bounded. But the level sets $\{z \mid g(z) \geq \alpha\}$, $\alpha \in R$, are all bounded if merely one of them is bounded and nonempty (see [7, Corollary 8.7.1]). In the present case we known that the level set

$$\{z \mid g(z) \geq \overline{\alpha}\} = [\text{set of all optimal solutions to (D)}]$$

is bounded and nonempty (Theorem 2). Therefore the set $\{z \in Z \mid G(z) \geq G(\overline{z}^1)\}$ is indeed bounded, and the sequence $\{\overline{z}^\nu\}$ is bounded as claimed.

We invoke now the fact that

(4.8) $\qquad \overline{x}^\nu \in \xi(c - E_\omega T_\omega^* \overline{z}_\omega^\nu)$ for all ν,

which is true by (2.5) because $(\overline{x}^\nu, \overline{z}^\nu)$ is a saddle point of L relative to $X \times Z^\nu$. In terms of the finite concave function φ we have

(4.9) $\qquad \xi(v) = \partial\varphi(v)$ for all $v \in R^n$.

Indeed, (1.4) defines φ as the conjugate of the closed proper concave function

$$\gamma(x) = \begin{cases} -\frac{1}{2}x \cdot Cx & \text{if } x \in X, \\ \infty & \text{if } x \notin X, \end{cases}$$

so $\partial\varphi(v)$ consists of the points x which minimize $v \cdot x - \gamma(x)$ over R^n (see [7, Theorem 23.5]). These are the points that make up the set $\xi(v)$ in (1.2). Thus

(4.10) $\qquad \overline{x}^\nu \in \partial\varphi(\overline{v}^\nu)$ for all ν, where $\overline{v}^\nu = c - E_\omega T_\omega^* \overline{z}_\omega^\nu$.

The sequence $\{\overline{v}^\nu\}$ is bounded, since $\{\overline{z}^\nu\}$ is. Moreover the multifunction $\partial\varphi$ is locally bounded: for every $\overline{v} \in R^n$ there is a $\delta > 0$ such that the set $\bigcup\{\partial\varphi(v) \mid |v - \overline{v}| \leq \delta\}$ is bounded (see [7, Corollary 24.5.1]). It follows by a simple compactness argument that $\partial\varphi$ carries bounded sets into bounded sets: if $V \subset R^n$ is bounded, then $\bigcup\{\partial\varphi(v) \mid v \in V\}$ is bounded. Taking $V = \{\overline{v}^\nu\}$, we conclude that the sequence $\{\overline{x}^\nu\}$ is bounded.

The argument establishing that $\{z^\nu\}$ is bounded is similar. We have $z_\omega^\nu \in \zeta_\omega(\overline{x}^\nu)$, where ζ_ω is the multifunction defined in (1.3). Since the sequence $\{\overline{x}^\nu\}$ is now known to be bounded, we need only show that ζ_ω is locally bounded at every $\overline{x}^\nu$ in order to conclude that each of the sequences $\{z_\omega^\nu\}$ is bounded and consequently that $\{z^\nu\}$ is bounded.

In terms of the convex function θ_ω defined in (1.5) we have

(4.11) $\qquad \zeta_\omega(x) = \partial\theta_\omega(h_\omega - T_\omega x)$ for all $x \in X$.

This holds because (1.5) expresses θ_ω as the conjugate of the closed proper convex function

$$F_\omega(z_\omega) = \begin{cases} \frac{1}{2}z_\omega \cdot H_\omega z_\omega & \text{if } z_\omega \in Z_\omega, \\ \infty & \text{if } z_\omega \notin Z_\omega, \end{cases}$$

The vectors $z_\omega \in \partial \theta_\omega(u)$ are therefore the ones that maximize $u \cdot z_\omega - f_\omega(z_\omega)$ (see [7, Theorem 23.5]). Our assumption that $\zeta_\omega(x)$ is nonempty and bounded for every $x \in X$ means that $\partial \theta_\omega(u)$ is nonempty and bounded for every u of the form $h_\omega - T_\omega x$ for some $x \in X$. Every such $u = h_\omega - T_\omega x$ therefore belongs to int(dom θ_ω) (cf. [7, Theorem 23.4]). It follows then that $\partial \theta_\omega$ is locally bounded at u (cf. [7, Corollary 24.5.1]). The mapping $x \to h_\omega - T_\omega x$ is continuous, so this implies ζ_ω is locally bounded at x for every $x \in X$, as we needed to prove.

The argument just given shows also that the convex function θ_ω is continuous $h_\omega - T_\omega x$ for every $x \in X$ (since θ_ω is continuous on int(dom θ_ω) [7, Theorem 10.1]). Therefore F is continuous on X by (1.6) and (2.6). We observed earlier in the proof that G is also continuous on Z. Of course X and Z, being convex polyhedra, are closed sets. Hence if $\varepsilon_\nu \to 0$, so that $F(\bar{x}^\nu) \to \bar{\alpha}$ and $G(\bar{z}^\nu) \to \bar{\alpha}$, any cluster points $\bar{x}^\infty$ of $\{\bar{x}^\nu\}$ and $\bar{z}^\infty$ of $\{\bar{z}^\nu\}$ must satisfy $F(\bar{x}^\infty) = \bar{\alpha} = G(\bar{z}^\infty)$ and be optimal solutions to (P) and (D).

We turn finally to the estimate (4.6). The saddle point condition on $(\bar{x}^\nu, \bar{z}^\nu)$ entails

$$\bar{x}^\nu \in \underset{x \in X}{\arg\min} \, L(x, \bar{z}^\nu).$$

Since X is a closed convex set and $L(x, \bar{z}^\nu)$ is a differentiable convex function of x, this condition implies that the vector $-\bar{w}^\nu = -\nabla_x L(\bar{x}^\nu, \bar{z}^\nu)$ belongs to the normal cone to X at $\bar{x}^\nu$ (cf. [7, Theorem 27.4]), which is exactly the assertion of (4.7). We have

$$(4.12) \qquad L(x, \bar{z}^\nu) = L(\bar{x}^\nu, \bar{z}^\nu) + \nabla_x L(\bar{x}^\nu, \bar{z}^\nu) \cdot (x - \bar{x}^\nu) + \tfrac{1}{2}(x - \bar{x}^\nu) \cdot C(x - \bar{x}^\nu)$$

$$= \bar{\alpha}_\nu = \bar{w}^\nu \cdot (x - \bar{x}^\nu) + \tfrac{1}{2} \| x - \bar{x}^\nu \|_C^2 \quad \text{for all} \quad x$$

from the quadratic nature of L, and also

$$L(x, \bar{z}^\nu) \le F(x) \quad \text{for all} \quad x \in X$$

by (2.2). For any optimal solution $\bar{x}$ to (P), then, we have

$$\bar{\alpha}_\nu + \bar{w}^\nu \cdot (\bar{x} - \bar{x}^\nu) + \tfrac{1}{2} \| \bar{x} - \bar{x}^\nu \|_C^2 \le F(\bar{x}) = \bar{\alpha}.$$

In terms of $\bar{\varepsilon}_\nu = \bar{\alpha} - \bar{\alpha}_\nu$, this can be written as the first inequality in (4.6). The rest of (4.6) then follows from (4.7), inasmuch as $\varepsilon_\nu = \alpha_\nu - \underline{\alpha}_\nu = \bar{\varepsilon}_\nu + \underline{\alpha}_\nu - \bar{\alpha} \ge \bar{\varepsilon}_\nu$. □

Theorem 3 focuses our attention on finding conditions that guarantee $\varepsilon_\nu \to 0$. Our first result in this direction makes no additional assumptions on the data in the problem and therefore serves as a baseline. It relies on an increasing sequence of polytopes in Step 4, however. The generalized cutting-plane rule in (3.32) is covered as a special case.

THEOREM 4. *If $Z^{\nu+1} \supset Z^\nu \cup \{z^\nu\}$ in Step 4 of the algorithm, then $\varepsilon_\nu \to 0$.*

PROOF. Let $\overline{\alpha}_\infty = \lim_\nu \overline{\alpha}_\nu$ and $\alpha_\infty = \lim \sup_\nu \alpha_\nu$. (The first limit exists because $\{\overline{\alpha}_\nu\}$ is nondecreasing in (4.5).) Since $\varepsilon_\nu = \alpha_\nu - \overline{\alpha}_\nu \geq 0$ for all ν, we need only demonstrate that $\alpha_\infty \leq \overline{\alpha}_\infty$. The sequences $\{\overline{x}^\nu\}$, $\{\overline{z}^\nu\}$, and $\{z^\nu\}$, are bounded by Theorem 3, so we can extract convergent subsequences with a common index set $N \subset \{1,2,...\}$ such that

$$\overline{x}^\nu \underset{\nu \in N}{\to} \overline{x}^\infty, \quad \overline{z}^\nu \underset{\nu \in N}{\to} \overline{z}^\infty, \quad z^\nu \underset{\nu \in N}{\to} z^\infty, \quad \alpha_\nu \underset{\nu \in N}{\to} \alpha_\infty.$$

Then since

$$\overline{\alpha}_\nu = L(\overline{x}^\nu, \overline{z}^\nu) \underset{\nu \in N}{\to} L(\overline{x}^\infty, \overline{z}^\infty), \quad \alpha_\nu = \overline{L(\overline{x}^\nu, z^\nu)} \underset{\nu \in N}{\to} L(\overline{x}^\infty, z^\infty),$$

we have $\overline{\alpha}_\infty = L(\overline{x}^\infty, \overline{z}^\infty)$ and $\alpha_\infty = L(\overline{x}^\infty, z^\infty)$. Our task now is to prove that $L(\overline{x}^\infty, z^\infty) \leq L(\overline{x}^\infty, \overline{z}^\infty)$.

From the saddle point condition on $(\overline{x}^\infty, \overline{z}^\infty)$ we have

$$L(\overline{x}^\nu, z) \leq L(\overline{x}^\nu, \overline{z}^\nu) \quad \text{for all } z \in Z^\nu.$$

Let $Z^\infty = \bigcup_{\nu=1}^\infty Z^\nu$. Since $Z^\nu \subset Z^{\nu+1} \subset \cdots$ we know that for any fixed $z \in Z^\infty$ the inequality $L(\overline{x}^\nu, z) \leq L(\overline{x}^\nu, \overline{z}^\nu)$ holds for all ν sufficiently high. Taking the limit as $\nu \to \infty$, $\nu \in N$, we obtain $L(\overline{x}^\infty, z) \leq L(\overline{x}^\infty, \overline{z}^\infty)$. This holds for arbitrary $z \in Z^\infty$, so

$$L(\overline{x}^\infty, z) \leq L(\overline{x}^\infty, \overline{z}^\infty) \quad \text{for all } z \in \text{cl } Z^\infty.$$

But z^∞ is one of the elements of $\text{cl } Z^\infty$, since $z_\nu \in Z^{\nu+1}$ for all ν. Therefore $L(\overline{x}^\infty, z^\infty) \leq L(\overline{x}^\infty, \overline{z}^\infty)$ in particular, and the proof is complete. □

Our main result comes next. It assures us that when C is positive definite, we do *not* have to keep increasing the size of the polytope Z^ν in order to have convergence. The number of elements used to generate Z^ν can be kept at whatever level seems adequate in maintaining a robust representation of F and G.

THEOREM 5. *Suppose the matrix C in (P) is positive definite. Then under the minimal requirement $Z^{\nu+1} \supset \{\overline{z}^\nu, z^\nu\}$ in Step 4 of the algorithm, one has $\varepsilon_\nu \to 0$ and also $\overline{x}^\nu \to \overline{x}$, where $\overline{x}$ is the unique optimal solution to (P).*

If in addition there exists $\rho \geq 0$ such that

(4.13)
$$z_\omega \cdot T_\omega C^{-1} T_\omega^* z_\omega \leq \rho z_\omega \cdot H_\omega z_\omega \quad \text{for all } z_\omega \in R^m, \omega \in \Omega,$$

(as is true in particular if every H_ω is positive definite), then in the estimate (4.6) one has

(4.14)
$$\overline{\varepsilon}_{\nu+1} \leq \tau \overline{\varepsilon}_\nu \quad \text{for } \nu = 1, 2, ...$$

where the factor $\tau \in [0,1)$ is given by

(4.15)
$$\tau = \begin{cases} \rho & \text{if } 0 \leq \rho \leq \frac{1}{2}, \\ 1 - \frac{1}{4}\rho^{-1} & \text{if } \rho \geq \frac{1}{2}. \end{cases}$$

Thus

(4.16) $\bar{\varepsilon}_{\nu+\mu} \leq \tau^{\mu}\bar{\varepsilon}_{\nu} \leq \tau^{\mu}\varepsilon_{\nu}$ *for* $\nu=1,2,\ldots,$ *and* $\mu=1,2,\ldots.$

Note that Theorem 5 asserts in (4.14) a linear rate of convergence of $\bar{\alpha}_{\nu}$ to $\bar{\alpha}$ with modulus τ, and the estimate (4.6) effectively translates this into a linear rate of convergence of $\bar{x}^{\nu}$ to $\bar{x}$ with modulus $\tau^{1/2}$. Indeed, from (4.6) and (4.16) we have

$$||\bar{x}-\bar{x}^{\nu+\mu}||_C \leq [2\tau^{\mu}\varepsilon_{\nu}]^{1/2} \text{ for } \nu=1,2,\ldots \text{ and } \mu=1,2,\ldots.$$

This is an unusual sort of result, because it applies not just to the tail of the sequence $\{\bar{x}^{\nu}\}$ but right from the beginning. Moreover the value of ε_{ν} is known in each iteration, and the value of $\tau\in[0,1)$ can be estimated in advance.

Theorem 5 makes no assertion about the convergence of $\{\bar{z}^{\nu}\}$ beyond the one in Theorem 3. Of course if there is a *unique* optimal solution $\bar{z}$ to (D), then by Theorem 3 we have $\bar{z}^{\nu}\to\bar{z}$ whenever $\varepsilon_{\nu}\to0$, as is the case here. In particular (D) has a unique optimal solution if the matrices H_{ω} are all positive definite.

The proof of Theorem 5 depends on further analysis of the dual objective function G. Essentially what we must provide is a lower estimate of G that ensures that the direction $z^{\nu}-\bar{z}^{\nu}$ determined in Step 2 of the algorithm is always a direction of ascent for G.

PROPOSITION 3. *Let*

(4.17) $f^{\nu}(w) = \max_{x\in X}\{(w-\bar{w}^{\nu})\cdot x - \bar{x}^{\nu}) -\frac{1}{2}(x-\bar{x}^{\nu})\cdot C(x-\bar{x}^{\nu})\}$ *for* $w\in R^n.$

Then f^{ν} *is a finite convex function on* R^n *with* $0 = f^{\nu}(0) \leq f^{\nu}(w)$ *for all* w, *and*

(4.18) $0\leq L(\bar{x}^{\nu},z)-G(z)=f^{\nu}(E_{\omega}T_{\omega}^*(z_{\omega}-\bar{z}_{\omega}^{\nu}))\leq E_{\omega}f^{\nu}(T_{\omega}^*(z_{\omega}-\bar{z}_{\omega}^{\nu}))$

for all $z\in Z.$

If C *is positive definite, then*

(4.19) $f^{\nu}(w)\leq\frac{1}{2}[(w-\bar{w}^{\nu})+s\bar{\omega}^{\nu}]\cdot C^{-1}[(w-\bar{w}^{\nu})+s\bar{w}^{\nu}]$ *for all* $s\geq0,$

so that in particular (for $s=1)$

(4.20) $G(z)\geq L(\bar{x}^{\nu},z)-\frac{1}{2}E_{\omega}\{(z_{\omega}-\bar{z}_{\omega}^{\nu})\cdot T_{\omega}^* C^{-1}T_{\omega}(z_{\omega}-\bar{z}_{\omega}^{\nu})\}$ *for all* $z\in Z.$

PROOF. First re-express f^{ν} in terms of the finite concave function φ in (1.4), so as to verify that f^{ν} is a finite convex function and that "max" rather than "sup" is appropriate in (4.17):

$$-f^{\nu}(w) = \min_{x\in X}\{(\bar{w}^{\nu}-w)\cdot(x-\bar{x}^{\nu}) +\frac{1}{2}(x-\bar{x}^{\nu})\cdot C(x-\bar{x}^{\nu})\}$$

$$= (w-\bar{w}^{\nu})\cdot\bar{x}^{\nu} +\frac{1}{2}\bar{x}^{\nu}\cdot C\bar{x}^{\nu} + \max_{x\in X}\{[\bar{w}^{\nu}-C\bar{x}^{\nu}-w]\cdot x +\frac{1}{2}x\cdot Cx\}$$

$$= (w - \overline{w}^\nu) \cdot \overline{x}^\nu + \tfrac{1}{2} \, \overline{x}^\nu \cdot C \overline{x}^\nu + \varphi(\overline{w}^\nu - C \overline{x}^\nu - w).$$

Clearly $f^\nu(w) \geq 0$ for all w, because $x = \overline{x}^\nu$ is one of the points considered in taking the maximum in (4.17). Furthermore

$$-f^\nu(0) = \min_{x \in X} \{ \overline{w}^\nu \cdot (x - \overline{x}^\nu) + \tfrac{1}{2} (x - \overline{x}^\nu) \cdot C (x - \overline{x}^\nu) \}.$$

Recalling the expansion (4.12) of $L(x, \overline{z}^\nu)$ around $\overline{x}^\nu$ and the fact that $\overline{x}^\nu$ minimizes $L(x, \overline{z}^\nu)$ over X (since $(\overline{x}^\nu, \overline{z}^\nu)$ is a saddle point of L on $X \times Z^\nu$), we see that $f^\nu(0) = 0$.

To get the equation in (4.18), from which the two inequalities in (4.18) immediately follow (the first because $f^\nu(w) \geq 0$ and the second by Jensen's inequality, because f^ν is convex), we look at the expansion

$$L(x, z) = L(\overline{x}^\nu, z) + \nabla_x L(\overline{x}^\nu, z) \cdot (x - \overline{x}^\nu) + \tfrac{1}{2} (x - \overline{x}^\nu) \cdot C (x - \overline{x}^\nu),$$

where

$$\nabla_x L(\overline{x}^\nu, z) = c + C \overline{x}^\nu - E_\omega T_\omega^* z_\omega = \overline{w}^\nu - E_\omega T_\omega^* (z_\omega - \overline{z}_\omega^\nu).$$

From this we calculate

$$L(\overline{x}^\nu, z) - G(z) = L(\overline{x}^\nu, z) - \min_{x \in X} L(x, z) = \max_{x \in X} \{ L(\overline{x}^\nu, z) - L(x, z) \}$$

$$= \max_{x \in X} \{ [E_\omega T_\omega^* (z_\omega - \overline{z}_\omega^\nu) - \overline{w}^\nu] \cdot (x - \overline{x}^\nu) - \tfrac{1}{2} (x - \overline{x}^\nu) \cdot C (x - \overline{x}^\nu) \}$$

$$= f^\nu(E_\omega T_\omega^* (z_\omega - \overline{z}_\omega^\nu)).$$

This establishes (4.18).

Finally we use property (4.7) in Theorem 1 to estimate for arbitrary $s \geq 0$:

$$f^\nu(w) \leq \sup_{x \in X} \{ [(w - \overline{w}^\nu) + s \, \overline{w}^\nu] \cdot (x - \overline{x}^\nu) - \tfrac{1}{2} (x - \overline{x}^\nu) \cdot C (x - \overline{x}^\nu) \}$$

$$\leq \sup_{x \in R^n} \{ [(w - \overline{w}^\nu) + s \, \overline{w}^\nu] \cdot (x - \overline{x}^\nu) - \tfrac{1}{2} (x - \overline{x}^\nu) \cdot C (x - \overline{x}^\nu) \}.$$

When C is positive definite, this last supremum equals the quadratic expression on the right side of (4.19). $\square$

PROOF OF THEOREM 5. Since $(\overline{x}^{\nu+1}, \overline{z}^{\nu+1})$ is a saddle point of L relative to $X \times Z^\nu$, we have

$$\overline{\alpha}_{\nu+1} = G(\overline{z}^{\nu+1}) = \max_{z \in Z^{\nu+1}} G(z).$$

But $Z^{\nu+1}$ includes the line segment joining $\overline{z}^\nu$ and z^ν. Therefore

$$(4.21) \qquad \overline{\alpha}_{\nu+1} \geq \max_{0 \leq t \leq 1} G(\overline{z}^\nu + t(z^\nu - \overline{z}^\nu)).$$

To see what this implies, we substitute $z = \overline{z}^\nu + t(z^\nu - \overline{z}^\nu)$ into the estimate (4.20) of Proposition 3 and make use of the fact that

$$(4.22) \qquad L(\overline{x}^\nu, \overline{z}^\nu + t(z^\nu - \overline{z}^\nu)) = L(\overline{x}^\nu, (1 - t)\overline{z}^\nu + t z^\nu)$$

$$\geq (1-t)L(\overline{x}^\nu,\overline{z}^\nu) + tL(\overline{x}^\nu,z^\nu) = (1-t)\overline{\alpha}_\nu + t\alpha_\nu = \overline{\alpha}_\nu + t\varepsilon_\nu$$

This yields

(4.23)
$$G(\overline{z}^\nu + t(z^\nu - \overline{z}^\nu)) \geq \overline{\alpha}_\nu + t\varepsilon_\varepsilon - \frac{1}{2}t^2\delta_\nu,$$

where

(4.24)
$$\delta_\nu := E_\omega\{(z_\omega^\nu - \overline{z}_\omega^\nu) \cdot TC^{-1}T^*(z_\omega^\nu - \overline{z}_\omega^\nu)\}.$$

Combining (4.23) with (4.21), we get

(4.25)
$$\overline{\alpha}_{\nu+1} \geq \overline{\alpha}_\nu + \sigma(\varepsilon_\nu, \delta_\nu),$$

where

(4.26)
$$\sigma(\varepsilon,\delta) := \max_{0 \leq t \leq 1}\{t\varepsilon - \frac{1}{2}t^2\delta\} = \begin{cases} \varepsilon - \frac{1}{2}\delta & \text{if} \quad 0 \leq \delta \leq \varepsilon, \\ \frac{1}{2}\varepsilon^2\delta^{-1} & \text{if} \quad \delta > \varepsilon. \end{cases}$$

Note that σ is a continuous function of $(\varepsilon,\delta) \in R_+^2$ with $\sigma(\varepsilon,\delta) = 0$ if $\varepsilon = 0$, but $\sigma(\varepsilon,\delta) > 0$ if $\varepsilon > 0$. The sequence $\{\overline{\alpha}_\nu\}$ is nondecreasing and bounded above by $\overline{\alpha}$ (cf. (4.5)), so $\sigma(\varepsilon_\nu,\delta_\nu) \to 0$. The sequence $\{\delta_\nu\}$ is bounded, because the sequences $\{\overline{z}^\nu\}$ and $\{z\}$ are bounded (Theorem 3). From the cited properties of σ, it follows then that $\varepsilon_\nu \to 0$. This implies $x^\nu \to \overline{x}$ by property (4.6) in Theorem 1.

We can also write (4.25) as

(4.27)
$$\overline{\varepsilon}_{\nu+1} \leq \overline{\varepsilon}_\nu - \sigma(\varepsilon_\nu,\delta_\nu).$$

Under the additional assumption in Theorem 5 that (4.13) holds, we have

(4.28)
$$\delta_\nu \leq \rho\beta_\nu, \quad \text{where} \quad \beta_\nu := E_\omega\{(z_\omega^\nu - \overline{z}_\omega^\nu) \cdot H_\omega(z_\omega^\nu - \overline{z}_\omega^\nu)\}.$$

Consider now the quadratic function

$$q(t) = L(\overline{x}^\nu, \overline{z}^\nu + t(z^\nu - \overline{z}^\nu)) \quad \text{for} \quad 0 \leq t \leq 1.$$

This has $q(0) = L(\overline{x}^\nu,\overline{z}^\nu) = \overline{\alpha}_\nu$, $q(1) = L(\overline{x}^\nu,z^\nu) = \alpha_\nu$, $q'' \equiv -\beta_\nu$, so q must be of the form

$$q(t) = (1-t)\overline{\alpha}_\nu + t\alpha_\nu + \frac{1}{2}t(1-t)\beta_\nu.$$

Moreover the maximum of $q(t)$ over $0 \leq t \leq 1$ is attained at $t = 1$, since the maximum of $L(\overline{x}^\nu, z)$ over $z \in Z$ is attained at $z = z^\nu$. Therefore

$$(1-t)\overline{\alpha}_\nu + t\alpha_\nu + \frac{1}{2}t(1-t)\beta_\nu \leq \alpha_\nu \text{ for } 0 \leq t \leq 1,$$

or in other words,

$$t(1-t)\beta_\nu \leq 2(1-t)(\alpha_\nu - \overline{\alpha}_\nu) = 2(1-t)\varepsilon_\nu \text{ for } 0 \leq t \leq 1.$$

This implies $\beta_\nu \leq 2\varepsilon_\nu$, and then (4.28) yields

(4.29)
$$\delta_\nu \leq 2\rho\varepsilon_\nu.$$

Formula (4.26) now gives us

$$\sigma(\varepsilon_\nu, \delta_\nu) \geq \sigma(\varepsilon_\nu, 2\rho\varepsilon_\nu) = \varepsilon_\nu \sigma(1, 2\rho) \geq \overline{\varepsilon}_\nu \sigma(1, 2\rho).$$

Substituting in (4.27) we get

(4.30) $$\overline{\varepsilon}_{\nu+1} \leq \overline{\varepsilon}_\nu - \overline{\varepsilon}_\nu \sigma(1, 2\rho) = [1 - \sigma(1, 2\rho)]\overline{\varepsilon}_\nu,$$

where

$$\sigma(1, 2\rho) = \begin{cases} 1 - \rho & \text{if} \quad 0 \leq \rho \leq \frac{1}{2} \\ \frac{1}{4}\rho & \text{if} \quad \rho \geq \frac{1}{2}. \end{cases}$$

The factor $1 - \sigma(1, 2\rho)$ is the number τ defined in (4.15), and (4.30) is thus the desired condition (4.14). []

REMARK. Proposition 3 provides additional information that could be used in the direction search and polytope modification steps in the algorithm. Inequality (4.18) asserts that

(4.31) $$L(\overline{x}^\nu, z) \geq G(z) \geq L(\overline{x}^\nu, z) - E_\omega f^\nu(T_\omega^*(z_\omega - \overline{z}_\omega^\nu))$$

$$\text{for all } z \in Z, \text{ with equality when } z = \overline{z}^\nu.$$

The vector z^ν maximizes $L(\overline{x}^\nu, z)$ over all $z \in Z$ and thus provides not only the needed value $L(\overline{x}^\nu, \overline{z}^\nu) = F(\overline{x}^\nu)$ but also a clue as to where we might try to move next in trying to improve on the current value $G(\overline{z}^\nu)$ of G. A further clue can be found by maximizing the right side of (4.31) over Z to get a vector $\hat{z}^\nu$. This is possible because the right side decomposes into separate terms for each ω. Indeed, the components $\hat{z}_\omega^\nu$ of $\hat{z}^\nu$ can be determined by

(4.32) $$\hat{z}_\omega^\nu \in \operatorname*{argmax}_{z_\omega \in Z_\omega} \{ f^\nu(T_\omega^*(z_\omega - \overline{z}_\omega^\nu)) + z_\omega \cdot [h_\omega - T_\omega \overline{x}^\nu] - \frac{1}{2} z_\omega \cdot H_\omega z_\omega \}.$$

In view of the form of f^ν in (4.17), this amounts to solving a special quadratic programming problem for each $\omega \in \Omega$.

If $\hat{z}^\nu$ is calculated in this way along with z^ν in Step 2, it can also be incorporated in the new polytope $Z^{\nu+1}$ in Step 4 in order to enrich the representation of G.

5. Adding Strongly Quadratic Terms.

The theoretical convergence properties of the finite generation algorithm are markedly superior when the quadratic forms that are involved are positive definite. But many problems lack this positive definiteness. Stochastic linear programming problems, for instance, have no quadratic terms at all. Such problems can be handled by a procedure which combines the finite generation algorithm with an augmented Lagrangian technique that introduces the desired property.

The technique in question was developed by Rockafellar [5] in a general context of minimax problems and variational inequalities. As applied to the present situation, it concerns the replacement of the saddle point problem for L on $X \times Z$ by a sequence of saddle point problems for augmented Lagrangians of the form

$$(5.1) \qquad L_*(x,z) = L(x,z) + \frac{\eta}{2}(x - \overline{x}_*^\mu) \cdot \overline{C}(x - \overline{x}_*^\mu)$$

$$- \frac{\eta}{2} E_\omega\{(z_\omega - \overline{z}_{*\omega}^\mu) \cdot \overline{H}_\omega(z_\omega - \overline{z}_{*\omega}^\mu)\} \text{ on } X \times Z \text{ for } \mu = 1,2,...$$

Here $\overline{C}$ and $\overline{H}_\omega$ are fixed positive definite matrices, η is a penalty parameter value that helps to control the rate of convergence, and $(\overline{x}_*^\mu, \overline{z}_*^\mu)$ is a current "estimate" for a saddle point of L itself on $X \times Z$, i.e. for an optimal solution pair for problems (P) and (D).

When the augmenting terms in L_* are expanded and combined with those in L, the expression (5.1) turns into

$$(5.2) \qquad L_*(x,z) = c_* \cdot x + \frac{1}{2} x \cdot C_* x$$

$$+ E_\omega\{z_\omega \cdot [h_{*\omega} - T_\omega x] - \frac{1}{2} z_\omega \cdot H_{*\omega} z_\omega\} + \text{const}.$$

where

$$(5.3) \qquad C_* = C + \eta\overline{C}, \quad H_{*\omega} = H_\omega + \eta\overline{H}_\omega,$$

$$(5.4) \qquad c_*^\mu = c - \eta\overline{C}\overline{x}_*^\mu, \quad h_{*\omega}^\mu = h_\omega - \eta\overline{H}_\omega z_{*\omega}^\mu$$

Note that the vectors c_*^μ and $h_{*\omega}^\mu$ giving the linear terms in L_μ depend on the μth solution estimates, but the matrices C_* and $H_{*\omega}$ giving the quadratic terms remain fixed as long as the value of η is not varied. Since $\eta > 0$, these matrices are positive definite. Therefore the saddle point problem for L_* on $X \times Z$ can be solved by the finite generation algorithm with an essentially linear rate of convergence (cf. Theorem 5).

We make use of this as follows.

MASTER ALGORITHM.

Step 0 (Initialization). Fix the matrices $\overline{C}$, $\overline{H}_\omega$, and the parameter value $\eta > 0$. Choose initial points $\overline{x}_*^1 \in X$ and $\overline{z}_*^1 \in Z$. Set $\mu = 1$.

Step 1 (Finite Generation Method). Use the finite generation algorithm to determine a saddle point $(\overline{x}_*, \overline{z}_*)$ of L_* on $X \times Z$.

Step 2 (Update). Set $(\overline{x}_*^{\mu+1}, \overline{z}_*^{\mu+1}) = (\overline{x}_*, \overline{z}_*)$. *Replace* μ *by* $\mu+1$ *and return to Step 1 (with the same value of* η *).*

The reader will note that Step 1 calls for an exact saddle point of L_*, whereas the finite generation method can only be expected to produce an approximate one. The theory in [5] on which the following theorem is based makes allowances for calculating only an approximate saddle point, but work still needs to be done on reconciling the stopping criterion in that theory with the possibilities available for the finite generation algorithm. For such reasons we content ourselves here with deriving from [5] only the simpler, qualitative results needed to put the approach in the right perspective.

THEOREM 6. *The sequences* $\{\overline{x}_*^{\mu}\}$ *and* $\{\overline{z}_*^{\mu}\}$ *converge to particular optimal solutions* $\overline{x}$ *and* $\overline{z}$ *to problems* (P) *and* (D), *respectively. If* $\overline{x}$ *and* $\overline{z}$ *are the unique optimal solutions to* (P) *and* (D), *then there is a number* $\beta(\eta) \in [0,1)$ *such that* $(\overline{x}_*^{\mu}, \overline{z}_*^{\mu})$ *converges to* $(\overline{x}, \overline{z})$ *at a linear rate with modulus* $\beta(\eta)$. *Moreover* $\beta(\eta) \to 0$ *as* $\eta \to 0$.

PROOF. We shall deduce this from [5, Theorem 5], which is a general result about the calculation of a saddle point of a convex-concave function on a product of Hilbert spaces. The Hilbert spaces in this case are R^n and $(R^m)^\Omega$ with different norms:

$$(5.5) \qquad \|x\|_* = [x \cdot \overline{C}x]^{1/2} \text{ for } x \in R^n,$$

$$\|z\|_* = [E_\omega\{z_\omega \cdot \overline{H}_\omega z_\omega\}]^{1/2} \text{ for } z \in (R^m)^\Omega.$$

The convex-concave function in question is

$$\overline{L}(x,z) = \begin{cases} L(x,z) & \text{if } x \in X \text{ and } z \in Z, \\ -\infty & \text{if } x \in X \text{ but } z \notin Z, \\ \infty & \text{if } x \notin X. \end{cases}$$

The saddle points of $\overline{L}$ on $R^n \times (R^m)^\Omega$ are the same as those of L on $X \times Z$. The problem of finding a saddle point of $\overline{L}_*$ on $X \times Z$ reduces to the one for

$$\overline{L}_*(x,z) = \overline{L}(x,z) + \frac{\eta}{2}\|x - \overline{x}_*^{\mu}\|_*^2 - \frac{\eta}{2}\|z - \overline{z}_*^{\mu}\|_*^2$$

on $R^n \times (R^m)^\Omega$.

The cited theorem of [5] concerns the sequence generated by taking $(\overline{x}_*^{\mu+1}, \overline{z}_*^{\mu+1})$ to be the saddle point of $\overline{L}_*$ on $R^n \times (R^m)^\Omega$, and this is the same as the sequence generated by our "master algorithm". The theorem ties in with others in [5] to give the convergence results we have claimed, provided that the following property holds when $(\overline{x}, \overline{z})$ is the *unique* saddle point of L on $X \times Z$: there exist $\gamma \geq 0$ and $\delta \geq 0$ such that the saddle points $(\tilde{x}, \tilde{z})$ of the perturbed Lagrangians of the form

$$\tilde{L}(x,y) = L(x,y) + \tilde{c} \cdot x + E_\omega\{\tilde{h}_\omega \cdot z_\omega\}$$

for various vectors $c \in R^n$ and $\tilde{h}_\omega \in R^m$ satisfy

$$\|(\tilde{x}, \tilde{z}) - \|_* \leq \gamma\|(\tilde{c}, \tilde{h})\|_* \text{ when } \|(\tilde{c}, \tilde{h})\| \leq \delta.$$

Here

$$\|(x,z)\|_* = [\|x\|_*^2 + \|z\|_*^2]^{1/2}.$$

This needed property does hold, because of the quadratic nature of our problem. The optimality conditions are all linear, so if they define a unique saddle point, the behavior of the saddle point with respect to perturbations will be Lipschitzian. []

We conclude by connecting the choice of the matrices $\overline{C}$ and $\overline{H}_\omega$ in (5.1) with the convergence rate of the finite generation algorithm in Step 1 of the master algorithm.

PROPOSITION 4. *Suppose $\overline{C}$ and $\overline{H}_\omega$ are selected so that for a certain $\overline{p} > 0$,*

(5.6) $$z_\omega \cdot [T_\omega \overline{C}^{-1} T_\omega^*] z_\omega \leq \overline{p}[z_\omega \cdot \overline{H}_\omega z_\omega] \quad \text{for all } z_\omega \in R^m.$$

Then the matrices C_* and $H_{*\omega}$ in (5.3) have

(5.7) $$z_\omega \cdot [T_\omega C_*^{-1} T_\omega^*] z_\omega \leq (\overline{p}/\eta^2)[z_\omega \cdot H_{*\omega} z_\omega] \quad \text{for all } z_\omega \in R^m,$$

so that when the finite generation algorithm is applied to finding a saddle point of L_*, the convergence results in Theorem 3 will be valid for

$$\rho = \overline{p}/\eta^2.$$

PROOF. Let us simplify notation by writing $A \leq B$ for positive definite symmetric matrices A and B to mean that $B - A$ is positive semidefinite. Since A and B can be diagonalized simultaneously, this relation can be interpreted also as a coordinatewise inequality on the corresponding vectors of eigenvalues. In this notation, our assumption (5.6) is that $T_\omega \overline{C}^{-1} T_\omega^* \leq \overline{p} \overline{H}_\omega$. Since $C_* = C + \eta \overline{C}$ we know $C_* \geq \eta \overline{C}$ and therefore $C_*^{-1} \leq \eta^{-1} \overline{C}^{-1}$. But also, from $H_{*\omega} = H_\omega + \eta \overline{H}_\omega$ we have $\eta \overline{H}_\omega \leq H_{*\omega}$, or in other words $\overline{H}_\omega \leq \eta^{-1} H_{*\omega}$. It follows that

$$T_\omega C_*^{-1} T_\omega^* \leq \eta^{-1} T_\omega \overline{C}^{-1} T_\omega^* \leq \eta^{-1} \overline{p} \overline{H}_\omega \leq \eta^{-2} \overline{p} H_{*\omega}$$

as claimed in (5.7). []

This result reveals a trade-off between the rates of linear convergence that can be achieved in the finite generation algorithm and in the master algorithm. The modulus $\beta(\eta)$ for the latter can be improved by making η smaller. But one cannot at the same time make ρ smaller, as would be desirable for the finite generation algorithm in the light of Theorem 5.

References

[1]   R.W. Cottle, "Symmetric dual quadratic programs," Quarterly of Applied Mathematics 21(1963) 237-243.

[2]   W.S. Dorn, "Duality in quadratic programming," Quarterly of Applied Mathematics 18(1960) 155-162.

[3]   M. Frank and P. Wolfe, "An algorithm for quadratic programming," *Naval Research Logistics Quarterly 3*(1956) 95-110.

[4]   H.W. Kuhn and A.W. Tucker, "Nonlinear programming," in: J. Neyman, ed., *Second Berkeley Symposium on Math. Statistics and Probability*, (Univ. of California Press, Berkeley, 1951).

[5]   R.T. Rockafellar, "Monotone operators and the proximal point algorithm," *SIAM Journal on Control and Optimization 14*(1976) 877-898.

[6]   R.T. Rockafellar, *Conjugate Duality and Optimization* (No. 16 in Conference Board of Math. Sci. Series, SIAM Publications, Philadelphia, 1974).

[7]   R.T. Rockafellar, *Convex Analysis* (Princeton University Press, Princeton, NJ, 1970).

[8]   R.T. Rockafellar and R.J.-B. Wets, "Stochastic convex programming: relatively complete recourse and induced feasibility," *SIAM Journal on Control and Optimization 14*(1976) 574-589.

[9]   R.T. Rockafellar and R.J.-B. Wets, "A dual solution procedure for quadratic stochastic programs with simple recourse," in: V. Pereyra and A. Reinoza, eds., *Numerical Methods* (Lecture Notes in Math. 1005, Springer-Verlag, Berlin, 1983) 252-265.

[10] R.J.-B. Wets, "Stochastic programming: solution techniques and approximation schemes," in: A. Bachem et al., eds., *Mathematical Programming: The State-of-the-Art* (Springer-Verlag, Berlin, 1983) 566-603.