Generalized Linear-Quadratic Problems of Deterministic and Stochastic Optimal Control in Discrete Time

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

SIAM J. Control and Optimization, 28:810–822, 1990.

Abstract. Two fundamental classes of problems in large-scale linear and quadratic programming are described. Multistage problems covering a wide variety of models in dynamic programming and stochastic programming are represented in a new way. Strong properties of duality are revealed which support the development of iterative approximate techniques of solution in terms of saddlepoints. Optimality conditions are derived in a form that emphasizes the possibilities of decomposition.

Keywords: discrete-time optimal control, dynamic programming, stochastic programming, large-scale linear-quadratic programming, intertemporal optimization, finite generation method.

* *Supported in part by grants from the National Science Foundation and the Air Force Office of Scientific Research.*

1. Introduction

The importance of linear and quadratic programming problems is well appreciated in finite-dimensional optimization. Such problems serve as mathematical models in their own right and as subproblems solved within the context of general numerical methods of nonlinear programming. In optimal control only a relatively small class of linear-quadratic problems has traditionally received much attention, however. A much more general class has recently been explored by Rockafellar [**1**] with the aim of opening up a wide domain for application of techniques of large-scale linear and quadratic programming, in particular the finite generation method of Rockafellar and Wets [**2**], [**3**], [**4**] that has been implemented in stochastic programming [**5**]. Central to this purpose is the development of flexible problem formulations for which there is a strong duality theory that represents optimal trajectories and controls in terms of saddlepoints of a "decomposable" Lagrangian.

In the present paper a discrete-time version of the deterministic models in [**1**] is investigated and corresponding results on optimality and duality are obtained. The formulations and results are then generalized to the stochastic case. The focus on discrete time is motivated by the computational possibilities already mentioned, so we do not hesitate to suppose also that the probability space for our stochastic version is discrete.

Our emphasis is on setting up a general framework for large-scale finite-dimensional linear-quadratic programming problems that reflect the special structure of optimal control. Besides being useful for numerical experimentation, such a framework may stimulate new applications, for instance in areas like operations research and resource systems management, where inequality constraints occur that jointly involve states and controls. Although the task of clarifying the relationship between finite and infinite-dimensional formulations is an important one, it is not the object of our efforts here.

In fact our discrete-time problems are more general than typical continuous-time problems in one respect: the dimensionality of the state and control vectors can vary with time. This feature is important in multistage modeling, where the decision structure in one period need not be the same as in another. The flexibility it provides allows us to show that a much wider class of problems is covered by our format than might at first be imagined.

2. Generalized Linear-Quadratic Programming.

The control problems that will be formulated are based on a concept of generalized linear-quadratic programming explained fully in Rockafellar [**1**]. A problem fits this concept if it can be expressed in the form

$$(\mathcal{P}) \qquad \text{minimize } f(u) = \sup_{v \in V} J(u,v) \text{ over all } u \in U,$$

where U and V are polyhedral convex sets in $\mathbb{R}^k$ and $\mathbb{R}^\ell$, and J is a quadratic convex-concave function on $U \times V$, namely

$$J(u,v) = p{\cdot}u + \tfrac{1}{2} u{\cdot}Pu + q{\cdot}v - \tfrac{1}{2} v{\cdot}Qv - v{\cdot}Du, \tag{2.1}$$

where P and Q are symmetric and positive *semi*definite (possibly 0—we do not exclude "linear" when we say "quadratic", as we try to underline by sometimes using the term "linear-quadratic"). The problem dual to $(\mathcal{P})$ is then

$$(\mathcal{Q}) \qquad \text{maximize } g(v) = \inf_{u \in U} J(u,v) \text{ over all } v \in V.$$

Here $f(u)$ could be ∞ and $g(v)$ could be $-\infty$. We regard u as a feasible solution to $(\mathcal{P})$ only if $u \in U$ and $f(u) < \infty$; likewise, we regard v as a feasible solution to $(\mathcal{Q})$ only if $v \in V$ and $g(v) > -\infty$.

The expression of problems $(\mathcal{P})$ and $(\mathcal{Q})$ is facilitated by the notation

$$\rho_{V,Q}(r) = \sup_{v \in V} \{ r{\cdot}v - \tfrac{1}{2} v{\cdot}Qv \} \text{ for } r \in \mathbb{R}^\ell, \tag{2.2}$$

$$\rho_{U,P}(s) = \sup_{u \in U} \{ s{\cdot}u - \tfrac{1}{2} u{\cdot}Pu \} \text{ for } s \in \mathbb{R}^k. \tag{2.3}$$

Thus $\rho_{V,Q}$ is a function on $\mathbb{R}^\ell$ determined by the specification of a polyhedral convex set $V \subset \mathbb{R}^\ell$ and a symmetric positive semidefinite matrix $Q \in \mathbb{R}^{\ell \times \ell}$. It is in general "piecewise linear-quadratic" in a sense made precise in [**1**], and it may take on the value ∞. There are many special cases deserving of mention, but for these too one should consult to [**1**]. Let it suffice to observe that when $0 \in V$, one has $\rho_{V,Q}(r) \geq 0$ for all r, $\rho_{V,Q}(0) = 0$. Then $\rho_{V,Q}(r)$ can be interpreted as an expression that "monitors deviations of r from 0". Similarly for $\rho_{U,P}$.

In this notation our general problems can be written as

$$(\mathcal{P}) \qquad \text{minimize } p{\cdot}u + \tfrac{1}{2} u{\cdot}Pu + \rho_{V,Q}(q - Du) \text{ over } u \in U,$$

$$(\mathcal{Q}) \qquad \text{maximize } q{\cdot}v - \tfrac{1}{2}v{\cdot}Qv - \rho_{U,P}(D^*v - p) \text{ over } v \in V$$

(where the asterisk * signals the transpose matrix). In $(\mathcal{P})$, therefore, one has the possibility of linear constraints represented by the condition $u \in U$, and also an objective term which "monitors deviations of Du from q". This may be a penalty term that is zero for some kinds of deviations but positive for others. For example, if $V = \mathbb{R}^\ell_+, Q = 0$, one has

$$\rho_{V,Q}(q - Du) = \begin{cases} 0 & \text{if } Du \geq q, \\ \infty & \text{if } Du \not\geq q, \end{cases} \tag{2.4}$$

so that the ρ term in $(\mathcal{P})$ is a "sharp" representation of the constraint $Du \geq q$. If at the same time one has $U = \mathbb{R}^k_+, P = 0$, then similarly

$$\rho_{U,P}(D^*v - p) = \begin{cases} 0 & \text{if } D^*v \leq p, \\ \infty & \text{if } D^*v \not\leq p. \end{cases} \tag{2.5}$$

In this case $(\mathcal{P})$ and $(\mathcal{Q})$ reduce to a canonical pair of linear programming problems in duality. See [**1**] for discussion of the rich possibilities that such ρ terms provide more generally in mathematical modeling.

The basic facts about the relationship between $(\mathcal{P})$ and $(\mathcal{Q})$ can be derived from the standard theory of linear and quadratic programming, specifically the duality theorem of Cottle [**6**] and the existence theorem of Frank and Wolfe [**7**].

Theorem 2.1 (Rockafellar and Wets [**3**, Theorem 2]). *If either $(\mathcal{P})$ or $(\mathcal{Q})$ has finite optimal value, or if both problems have feasible solutions, then both optimal values are finite and equal, and both problems have optimal solutions. In this case a pair $(\overline{u}, \overline{v})$ is a saddlepoint of $J(u, v)$ relative to $u \in U$ and $v \in V$ if and only if $\overline{u}$ is an optimal solution to $(\mathcal{P})$ and $\overline{v}$ is an optimal solution to $(\mathcal{Q})$.*

3. Deterministic Control Model.

We want now to formulate problems in this vein that belong to optimal control. The dynamical system we consider takes the form

$$
\begin{aligned}
x_\tau &= A_\tau x_{\tau-1} + B_\tau u_\tau + b_\tau \text{ for } \tau = 1, \dots, T, \\
x_0 &= B_0 u_0 + b_0, \qquad \text{where } u_\tau \in U_\tau \text{ for } \tau = 0, 1, \dots, T.
\end{aligned} \tag{3.1}
$$

The vectors $u_\tau \in \mathbb{R}^{k_\tau}$ are controls, and the vectors $x_\tau \in \mathbb{R}^{n_\tau}$ are states (observe that dimensions can vary with τ). We write $u = (u_0, u_1, \dots, u_T)$ and $x = (x_0, x_1, ..., x_T)$. Thus x is uniquely determined by u, and the transformation $u \mapsto x$ is affine. Note that u_0 serves as a supplementary parameter vector more than as a control vector in the usual dynamical sense.

The sets $U_\tau \subset \mathbb{R}^{k_\tau}$ are assumed to be polyhedral convex (nonempty). The matrices A_τ, B_τ and vectors b_τ are of appropriate dimension:

$$A_\tau \in \mathbb{R}^{n_\tau \times n_{\tau-1}}, \quad B_\tau \in \mathbb{R}^{n_\tau \times k_\tau}, \quad b_\tau \in \mathbb{R}^{n_\tau}.$$

(By taking $k_0 = 0$, one could eliminate u_0 from (3.1) and have $x_0 = b_0$.)

Our deterministic control problem is:

$$
\begin{aligned}
&\text{minimize subject to (3.1) the expression } f(u) = \\
&\sum_{\tau=0}^{T} [p_\tau{\cdot}u_\tau + \tfrac{1}{2} u_\tau{\cdot}P_\tau u_\tau - c_{\tau+1}{\cdot}x_\tau] \\
&+ \sum_{\tau=1}^{T} \rho_{V_\tau, Q_\tau}(q_\tau - C_\tau x_{\tau-1} - D_\tau u_\tau) + \rho_{V_{T+1}, Q_{T+1}}(q_{T+1} - C_{T+1} x_T).
\end{aligned} \tag{$\mathcal{P}_{\text{det}}$}
$$

Here V_τ is a polyhedral convex set (nonempty) in $\mathbb{R}^{\ell_\tau}$, and the matrices P_τ and Q_τ are symmetric and positive semidefinite. One has

$$P_\tau \in \mathbb{R}^{k_\tau \times k_\tau}, \quad Q_\tau \in \mathbb{R}^{\ell_\tau \times \ell_\tau}, \quad p_\tau \in \mathbb{R}^{k_\tau}, \quad q_\tau \in \mathbb{R}^{\ell_\tau},$$

$$c_\tau \in \mathbb{R}^{n_{\tau-1}}, \quad C_\tau \in \mathbb{R}^{\ell_\tau \times n_{\tau-1}}, \quad D_\tau \in \mathbb{R}^{\ell_\tau \times k_\tau}.$$

In this notation the elements A_τ and D_τ are defined only for $\tau = 1, \dots, T$, but $B_\tau, b_\tau, P_\tau, p_\tau$, are defined for $\tau = 0, 1, \dots, T$ and $C_\tau, c_\tau, Q_\tau, q_\tau$ for $\tau = 1, \dots, T, T+1$.

For the problem that will turn out to be dual to $(\mathcal{P}_{\text{det}})$, the dynamical system goes backward in time:

$$
\begin{aligned}
y_\tau &= A_\tau^* y_{\tau+1} + C_\tau^* v_\tau + c_\tau \text{ for } \tau = 1, \dots, T, \\
y_{T+1} &= C_{T+1}^* v_{T+1} + c_{T+1}, \qquad \text{where } v_\tau \in V_\tau \text{ for } \tau = 1, \dots, T, T+1.
\end{aligned} \tag{3.2}
$$

The vectors $v_\tau \in \mathbb{R}^{\ell_\tau}$ are the dual controls, and the vectors $y_\tau \in \mathbb{R}^{n_{\tau-1}}$ are the dual states. We write

$$v = (v_1, \ldots, v_T, v_{T+1}) \text{ and } y = (y_1, \ldots, y_T, y_{T+1}).$$

The dual problem then is

$$(\mathcal{Q}_{\text{det}}) \qquad \begin{aligned} &\text{maximize subject to (3.2) the expression } g(v) = \\ &\sum_{\tau=1}^{T+1} [q_\tau{\cdot}v_\tau - \tfrac{1}{2} v_\tau{\cdot}Q_\tau v_\tau - b_{\tau-1}{\cdot}y_\tau] \\ &- \sum_{\tau=1}^{T} \rho_{U_\tau,P_\tau}(B_\tau^* y_{\tau+1} - D_\tau^* v_\tau - p_\tau) - \rho_{U_0,P_0}(B_0^* y_1 - p_0). \end{aligned}$$

In this formula y is the trajectory uniquely determined from v by (3.2).

Proposition 3.1. *Suppose x corresponds to u by (3.1), and y to v by (3.2). Then*

$$\sum_{\tau=0}^{T} y_{\tau+1}{\cdot}[B_\tau u_\tau + b_\tau] = \sum_{\tau=1}^{T+1} x_{\tau-1}{\cdot}[C_\tau^* v_\tau + c_\tau]. \tag{3.3}$$

Proof. In view of the relations (3.1) the left side of (3.3) can be written as

$$\begin{aligned} &y_1{\cdot}x_0 + \sum_{\tau=1}^{T} y_{\tau+1}[x_\tau - A_\tau x_{\tau-1}] \\ &\qquad = y_1{\cdot}x_0 + y_2{\cdot}x_1 + \cdots + y_{T+1}{\cdot}x_T - \sum_{\tau=1}^{T} x_{\tau-1}{\cdot}A_\tau^* y_{\tau+1}. \end{aligned}$$

Likewise from (3.2) the right side becomes

$$\begin{aligned} &x_T{\cdot}y_{T+1} + \sum_{\tau=1}^{T} x_{\tau-1}{\cdot}[y_\tau - A_\tau^* y_{\tau+1}] \\ &\qquad = y_1{\cdot}x_0 + y_2{\cdot}x_1 + \cdots + y_{T+1}{\cdot}x_T - \sum_{\tau=1}^{T} x_{\tau-1}{\cdot}A_\tau^* y_{\tau+1}. \end{aligned}$$

Thus the two sides are equal, as claimed. □

Proposition 3.2. *Let $U = U_0 \times \cdots \times U_T$ and $V = V_1 \times \cdots \times V_{T+1}$, and for $u \in U$ and $v \in V$ define*

$$\begin{aligned} J(u,v) = &\sum_{\tau=0}^{T} (p_\tau{\cdot}u_\tau + \tfrac{1}{2} u_\tau{\cdot}P_\tau u_\tau) + \sum_{\tau=1}^{T+1} (q_\tau{\cdot}v_\tau - \tfrac{1}{2} v_\tau{\cdot}Q_\tau v_\tau) \\ &- \sum_{\tau=1}^{T} v_\tau{\cdot}D_\tau u_\tau - [u,v], \end{aligned} \tag{3.4}$$

where $[u,v]$ denotes the common value of the expression in (3.3).

Then U and V are polyhedral convex sets, and J is a quadratic convex-concave function.

Proof. This is immediate from our assumptions and the fact the expression $[u,v]$ is affine in u and v separately. □

Theorem 3.3. *The deterministic optimal control problems $(\mathcal{P}_{\rm det})$ and $(\mathcal{Q}_{\rm det})$ are the primal and dual problems of generalized linear-quadratic programming associated with the U, V, and J in Proposition 3.2. In particular, the assertions of Theorem 2.1 are valid for $(\mathcal{P}_{\rm det})$ and $(\mathcal{Q}_{\rm det})$.*

Proof. We need only show that the expressions $f(u)$ and $g(v)$ in $(\mathcal{P}_{\rm det})$ and $(\mathcal{Q}_{\rm det})$ arise according to the pattern in the general problems $(\mathcal{P})$ and $(\mathcal{Q})$ of §1. First using for $[u,v]$ in (3.4) the right hand expression in (3.3), we write

$$
\begin{aligned}
J(u,v) = &\sum_{\tau=0}^{T}(p_\tau{\cdot}u_\tau + \tfrac{1}{2}u_\tau{\cdot}P_\tau u_\tau) - \sum_{\tau=1}^{T+1} c_\tau{\cdot}x_{\tau-1} \\
&+\sum_{\tau=1}^{T}([q_\tau - C_\tau x_{\tau-1} - D_\tau u_\tau]{\cdot}v_\tau - \tfrac{1}{2}v_\tau{\cdot}Q_\tau v_\tau) \\
&\qquad + ([q_{T+1} - C_{T+1}x_T]{\cdot}v_{T+1} - \tfrac{1}{2}v_{T+1}{\cdot}Q_{T+1}v_{T+1})
\end{aligned}
\tag{3.5}
$$

The maximization of this over all $v \in V$ reduces to a separate maximization with respect to each of the components v_τ of v. Since by definition

$$
\sup_{v_\tau \in V_\tau} \{[q_\tau - C_\tau x_{\tau-1} - D_\tau u_\tau]{\cdot}v_\tau - \tfrac{1}{2}v_\tau{\cdot}Q_\tau v_\tau\} = \rho_{V_\tau,Q_\tau}(q_\tau - C_\tau x_{\tau-1} - D_\tau u_\tau)
$$

and

$$
\sup_{v_{T+1}\in V_{T+1}} \{[q_{T+1} - C_{T+1}x_T{\cdot}v_{T+1} - \tfrac{1}{2}v_{T+1}{\cdot}Q_{T+1}v_{T+1}\} = \rho_{V_{T+1},Q_{T+1}}(q_{T+1} - C_{T+1}x_T),
$$

we conclude that $\sup_{v\in V} J(u,v)$ is the $f(u)$ in $(\mathcal{P}_{\rm det})$.

Next using for $[u,v]$ the left hand expression in (3.3), we write

$$
\begin{aligned}
J(u,v) = &\sum_{\tau=1}^{T+1}(q_\tau{\cdot}v_\tau - \tfrac{1}{2}v_t{\cdot}Q_\tau v_\tau) - \sum_{\tau=0}^{T} b_\tau{\cdot}y_{\tau+1} \\
&\quad - \sum_{\tau=1}^{T}([B_\tau^* y_{\tau+1} + D_\tau^* v_\tau - p_\tau]{\cdot}u_\tau - \tfrac{1}{2}u_\tau{\cdot}P_\tau u_\tau) \\
&\qquad - ([B_0^* y_1 - p_0]{\cdot}u_0 - \tfrac{1}{2}u_0{\cdot}P_0 u_0).
\end{aligned}
\tag{3.6}
$$

The minimization of this over all $u \in U$ reduces similarly to a separate minimization with respect to each of the components u_τ. We know that

$$\sup_{u_\tau \in U_\tau} \{[B_\tau^* y_{\tau+1} + D_\tau^* v_\tau - p_\tau]{\cdot}u_\tau - \tfrac{1}{2} u_\tau{\cdot}P_\tau u_\tau\} = \rho u_\tau, P_\tau(B_\tau^* y_{\tau+1} + D_\tau^* v_\tau - p_\tau)$$

and

$$\sup_{u_0 \in U_0} \{[B_0^* y_1 - p_0]{\cdot}u_0 - \tfrac{1}{2} u_0{\cdot}P_0 u_0\} = \rho_{u_0,P_0}(B_0^* y_1 - p_0).$$

We conclude that $\inf_{u \in U} J(u,v)$ is the $g(v)$ in $(\mathcal{Q}_{\rm det})$. □

The proof of Theorem 3.3 reveals an important simplifying feature of our minimax representation of $(\mathcal{P}_{\rm det})$ and $(\mathcal{Q}_{\rm det})$. We state it as follows.

Theorem 3.4. *For the U, V, and J in Theorem 3.3 one has the following decomposability properties for separate minimization in u or maximization in v. Here $\overline{u}$ and $\overline{v}$ are elements of U and V, and $\overline{x}$ and $\overline{y}$ the corresponding trajectories.*

(a) $\tilde{u} \in \operatorname*{argmin}_{u \in U} J(u, \overline{v})$ *if and only if*

$$\begin{aligned}\tilde{u}_\tau &\in \partial\rho_{U_\tau,P_\tau}(B_\tau^* \overline{y}_{\tau+1} + D_\tau^* \overline{v}_\tau - p_\tau)\\ &= \operatorname*{argmax}_{u_\tau \in U_\tau}\{[B_\tau^* \overline{y}_{\tau+1} + D_\tau^* \overline{v}_\tau - p_\tau]{\cdot}u_\tau - \tfrac{1}{2}u_\tau{\cdot}P_\tau u_\tau]\}\end{aligned}$$

for $\tau = 1, \ldots, T$, *and*

$$\begin{aligned}\tilde{u}_0 &\in \partial\rho_{U_0,P_0}(B_0^* \overline{y}_1 - p_0)\\ &= \operatorname*{argmax}_{u_0 \in U_0}\{[B_0^* \overline{y}_1 - p_0] - \tfrac{1}{2}u_0{\cdot}P_0 u_0\}.\end{aligned}$$

(b) $\tilde{v} \in \operatorname*{argmax}_{v \in V} J(\overline{u}, v)$ *if and only if*

$$\begin{aligned}\tilde{v}_\tau &\in \partial\rho_{v_\tau,Q_\tau}(q_\tau - C_\tau \overline{x}_{\tau-1} - D_\tau \overline{u}_\tau)\\ &= \operatorname*{argmax}_{v_\tau \in V_\tau}\{[q_\tau - C_\tau \overline{x}_{\tau-1} - D_\tau \overline{u}_\tau]{\cdot}v_\tau - \tfrac{1}{2}v_\tau{\cdot}Q_\tau v_\tau\}\end{aligned}$$

for $\tau = 1, \ldots, T$, *and*

$$\begin{aligned}\tilde{v}_{T+1} &\in \partial\rho_{V_{T+1},Q_{T+1}}(q_{T+1} - C_{T+1}\overline{x}_T)\\ &= \operatorname*{argmax}_{v_{T+1} \in V_{T+1}}\{[q_{T+1} - C_{T+1}\overline{x}_t]{\cdot}v_{T+1} - \tfrac{1}{2}v_{T+1}{\cdot}Q_{T+1}v_{T+1}\}.\end{aligned}$$

Proof. The formulas in terms of "argmax" are justified by the calculations in the proof of Theorem 3.3. The question that remains is whether the "argmax" sets are truly the same as the indicated subgradient sets. This is answered by the observation that in the notation (2.2) one has $\rho_{V,Q} = \theta_{V,Q}^*$ (convex conjugate), where

$$\theta_{V,Q}(v) = \begin{cases} \tfrac{1}{2}v{\cdot}Qv & \text{if } v \in V,\\ \infty & \text{if } v \notin V.\end{cases} \tag{3.7}$$

Inasmuch as $\theta_{V,Q}$ is a closed proper convex function, one also has $\theta_{V,Q} = \rho^*_{V,Q}$ and

$$\partial\rho_{V,Q}(r) = \operatorname*{argmax}_{v\in\mathbb{R}^\ell}\{r\cdot v - \theta_{V,Q}(v)\}$$

by the basic rules of convex analysis [**8**, Theorem 12.2]. When this is applied to the pairs V_τ, Q_τ, and U_τ, P_τ, in place of V, Q, we reach our desired conclusion. □

The significance of the formulas in Theorem 3.4 lies in their potential use in iterative methods for solving $(\mathcal{P}_{\rm det})$ and $(\mathcal{Q}_{\rm det})$ when the dimensions

$$k = \sum_{\tau=0}^{T} k_\tau \text{ and } \ell = \sum_{\tau=1}^{T+1} \ell_\tau \tag{3.9}$$

of the vectors $u = (u_0, u_1, \ldots, u_T)$ and $v = (v_1, \ldots, v_T, v_{T+1})$ are large. The dimensions may be expected to be large if T is large, as of course would happen in particular in taking $(\mathcal{P}_{\rm det})$ and $(\mathcal{Q}_{\rm det})$ to be discrete-time approximations to continuous-time control problems such as the ones studied in [**1**]. In the presence of high dimensions, it may be impossible or inexpedient to solve $(\mathcal{P}_{\rm det})$ and $(\mathcal{Q}_{\rm det})$ directly by reducing them to ordinary quadratic programming problems in duality and applying a typical finitely-terminating quadratic programming code (as would be possible in principle in a manner explained in Rockafellar and Wets [**3**,§2]).

An alternative approach in that case is the exploration of methods that determine approximate solutions to $(\mathcal{P}_{\rm det})$ and $(\mathcal{Q}_{\rm det})$ by calculating a sequence of approximate saddlepoints $(\overline{u}^\nu, \overline{v}^\nu)$ of J on $U \times V$ for $\nu = 1, 2, \ldots$, as suggested by the characterization of optimality in Theorem 3.4. In any such method the ability to calculate

$$f(\overline{u}^\nu) = \max_{v\in V} J(\overline{u}^\nu, v) \text{ and } \tilde{v}^\nu \in \operatorname*{argmax}_{v\in V} J(\overline{u}^\nu, v) \tag{3.10}$$

as well as

$$g(\overline{v}^\nu) = \min_{u\in U} J(u, \overline{v}^\nu) \text{ and } \tilde{u}^\nu \in \operatorname*{argmin}_{u\in U} J(u, \overline{v}^\nu) \tag{3.11}$$

is crucial in producing primal and dual bounds that tell how far $\overline{u}^\nu$ and $\overline{v}^\nu$ are from optimality and as input to possible schemes for updating $(\overline{u}^\nu, \overline{v}^\nu)$ to $(\overline{u}^{\nu+1}, \overline{v}^{\nu+1})$. Theorem 3.4 says that the calculations in (3.10) and (3.11) can feasibly be carried out in terms of solving a collection of low-dimensional quadratic programming subproblems indexed by τ. Moreover these subproblems can even be solved in "closed form", i.e. without applying a quadratic programming code, if the functions ρ_{V_τ,Q_τ} and ρ_{U_τ,P_τ} have sufficiently simple expressions that allow the use of subgradient formulas directly.

The subgradient formulas are readily usable, for example, in the completely decomposable case where U_τ and V_τ are boxes (products of closed intervals, e.g. orthants) and P_τ and Q_τ are diagonal. Indeed, if P_τ and Q_τ are nonsingular the subgradients reduce to gradients given by very elementary expressions.

Theorem 3.5. *Consider a control pair $\overline{u}, \overline{v}$, and the corresponding trajectories $\overline{x}$ and $\overline{y}$ determined by (3.1) and (3.2). Define*

$$\text{(3.12)}\quad \overline{p}_\tau = p_\tau - B_\tau^* \overline{y}_{\tau+1} \text{ for } \tau = 0, 1, \ldots, T, \text{ and } \overline{q}_\tau = q_\tau - C_\tau \overline{x}_{T-1} \text{ for } \tau = 1, \ldots, T, T+1.$$

Let $(\overline{\mathcal{P}}_\tau)$ and $(\overline{\mathcal{Q}}_\tau)$ for $\tau = 1, \ldots, T$ denote the primal and dual problems of generalized linear-quadratic programming associated with

$$\text{(3.13)}\qquad J_\tau(u_\tau, v_\tau) = \overline{p}_\tau \cdot u_\tau + \tfrac{1}{2} u_\tau \cdot P_\tau u_\tau + \overline{q}_\tau \cdot v_\tau - \tfrac{1}{2} v_\tau \cdot Q_\tau v_\tau - v_\tau \cdot D_\tau u_\tau$$

on $U_\tau \times V_\tau$, namely,

$$(\overline{\mathcal{P}}_\tau)\qquad \text{minimize } \overline{p}_\tau \cdot u_\tau + \tfrac{1}{2} u_\tau \cdot P_\tau u_\tau + \rho_{V_\tau, Q_\tau}(\overline{q}_\tau - D_\tau u_\tau) \text{ over } u_\tau \in U_\tau,$$

$$(\overline{\mathcal{Q}}_\tau)\qquad \text{maximize } \overline{q}_\tau \cdot v_\tau - \tfrac{1}{2} v_\tau \cdot Q_\tau v_\tau - \rho_{U_\tau, P_\tau}(D_\tau^* v_\tau - \overline{p}_\tau) \text{ over } v_\tau \in V_\tau,$$

and consider also the problems

$$(\overline{\mathcal{P}}_0)\qquad \text{minimize } \overline{p}_0 \cdot u_0 + \tfrac{1}{2} u_0 \cdot P_0 u_0 \text{ over } u_0 \in U_0,$$

$$(\overline{\mathcal{Q}}_{T+1})\qquad \text{maximize } \overline{q}_{T+1} \cdot v_{T+1} - \tfrac{1}{2} v_{T+1} \cdot Q_{T+1} v_{T+1} \text{ over } v_{T+1} \in U_{T+1}.$$

Then a necessary and sufficient condition for $\overline{u}$ and $\overline{v}$ to be optimal solutions to the control problems $(\mathcal{P}_{\rm det})$ and $(\mathcal{Q}_{\rm det})$, respectively, is that $\overline{u}_\tau$ should be an optimal solution to the subproblem $(\overline{\mathcal{P}}_\tau)$ for $\tau = 0, 1, \ldots, T$, and $\overline{v}_\tau$ should be an optimal solution to the subproblem $(\overline{\mathcal{Q}}_\tau)$ for $\tau = 1, \ldots, T, T+1$.

Proof. We know from Theorem 3.3 that a necessary and sufficient condition for the optimality of $\overline{u}$ and $\overline{v}$ in $(\mathcal{P}_{\rm det})$ and $(\mathcal{Q}_{\rm det})$ is the saddlepoint relation

$$\overline{u} \in \operatorname*{argmin}_{u \in U} J(u, \overline{v}) \text{ and } \overline{v} \in \operatorname*{argmax}_{v \in V} J(\overline{u}, v).$$

Furthermore, this reduces to having the argmax conditions in Theorem 3.4 hold for $\tilde{u} = \overline{u}$ and $\tilde{v} = \overline{v}$. These conditions in turn are equivalent to

$$\begin{aligned} \overline{u}_\tau &\in \operatorname*{argmin}_{u_\tau \in U_\tau} J_\tau(u_\tau, \overline{v}_\tau) \text{ for } \tau = 1, \ldots, T, \\ \overline{u}_0 &\in \operatorname*{argmin}_{u_0 \in U_0} \{ \overline{p}_0 \cdot u_0 + \tfrac{1}{2} u_0 \cdot P_0 u_0 \}, \end{aligned}$$

and

$$\begin{aligned} \overline{v}_\tau &\in \operatorname*{argmax}_{v_\tau \in V_\tau} J_\tau(\overline{u}_\tau, v_\tau) \text{ for } \tau = 1, \ldots, T, \\ \overline{v}_{T+1} &\in \operatorname*{argmax}_{v_{T+1} \in V_{T+1}} \{ \overline{q}_{T+1} \cdot v_{T+1} - \tfrac{1}{2} v_{T+1} \cdot Q_{T+1} v_{T+1} \}. \end{aligned}$$

The latter mean that $\overline{u}_0$ is optimal for $(\mathcal{P}_0)$, $\overline{v}_{T+1}$ is optimal for $(\mathcal{Q}_{T+1})$, and $(\overline{u}_\tau, \overline{v}_\tau)$ is a saddlepoint of $J_\tau(u_\tau, v_\tau)$ relative to $u_\tau \in U_\tau$ and $v_\tau \in V_\tau$ for $\tau = 1, \ldots, T$. This saddlepoint condition is equivalent by Theorem 2.1 to $\overline{u}_\tau$ and $\overline{v}_\tau$ being optimal solutions to the primal and dual subproblems $(\overline{\mathcal{P}}_\tau)$ and $(\overline{\mathcal{Q}}_\tau)$. □

Optimality conditions of the kind in Theorem 3.5 were developed for continuous-time problems in Rockafellar [**1**]. They resemble conditions first detected in a special setting known as "continuous linear programming" by Grinold [**9**].

Besides being of interest in the study of what optimality might mean in a particular application modeled directly in terms of $(\mathcal{P}_{\text{det}})$ and $(\mathcal{Q}_{\text{det}})$, the conditions in Theorem 3.5, like those in Theorem 3.4, have import for computations. Having arrived at a control pair $(\overline{u}^\nu, \overline{v}^\nu)$ and associated trajectories $(\overline{x}^\nu, \overline{y}^\nu)$ in some iteration ν of a numerical method, one can construct a new pair $(u^\nu, v^\nu) \in U \times V$ by taking u^ν_τ to be an optimal solution to $(\overline{\mathcal{P}}^\nu_\tau)$ for $\tau = 0, 1, \ldots, T$ and v^ν_τ an optimal solution to $(\overline{\mathcal{Q}}^\nu_\tau)$ for $\tau = 1, \ldots, T, T+1$, where $(\overline{\mathcal{P}}^\nu_\tau)$ and $(\overline{\mathcal{Q}}^\nu_\tau)$ are the subproblems corresponding to $\overline{u}^\nu$ and $\overline{v}^\nu$ in the sense of Theorem 3.5. Then u^ν and v^ν generate new trajectories x^ν and y^ν that may be compared with $\overline{x}^\nu$ and $\overline{y}^\nu$, and for so forth. This procedure, like the one described after Theorem 3.4, provides another tool that might be used constructively in the generation of a sequence of approximate saddlepoints.

4. Stochastic Control Model.

The probability space we work with in this paper is simply a finite set Ω, for reasons given in §1. The probability associated with an element $\omega \in \Omega$ is $\pi_\omega \geq 0$; one has $\sum_{\omega \in \Omega} \pi_\omega = 1$. The vectors, matrices and sets introduced in the formulation of our deterministic problems persist notationally in the stochastic problems, but all are now treated as (potentially) random variables. Thus, for example, p_τ now denotes a mapping $\omega \mapsto p_{\omega\tau} \in \mathbb{R}^{k_\tau}$ rather than necessarily just a single vector. Likewise P_τ is a matrix-valued mapping $\omega \mapsto P_{\omega\tau}$, and U_τ is a set-valued mapping $\omega \mapsto U_{\omega\tau}$. In line with our earlier assumptions, we suppose that $P_{\omega\tau}$ and $Q_{\omega\tau}$ are *positive semidefinite* (symmetric), and $U_{\omega\tau}$ and $V_{\omega\tau}$ are *polyhedral convex* (nonempty). The expectation of a random variable such as p_τ is

$$E\{p_\tau\} = E_\omega\{p_{\omega\tau}\} := \sum_{\omega \in \Omega} \pi_\omega p_{\omega\tau}.$$

The information available to the decision-making process at time τ is modeled by the specification of a (finite) field $\mathcal{G}_\tau$ of subsets of Ω for $\tau = 0, 1, \ldots, T, T+1$. The fields $\mathcal{G}_\tau$ may differ from the complete information fields $\mathcal{F}_\tau$, and no particular relation between them is presupposed, although the case where the $\mathcal{G}_\tau$'s are increasing with $\mathcal{G}_\tau$ contained in $\mathcal{F}_\tau$ is, for instance, an important one. More will be said about this after the statement of our primal and dual problems. We assume that

$$U_\tau, V_\tau, p_\tau, P_\tau, q_\tau, Q_\tau, \text{ and } D_\tau \text{ are } \mathcal{G}_\tau\text{-measurable,} \tag{4.1}$$

but in general do *not* place this restriction on A_τ, B_τ, C_τ, b_τ or c_τ. Trivially the latter are measurable with respect to the underlying field $\mathcal{F}$ of complete information, comprised here of *all* the subsets of Ω.

Because $\mathcal{G}_\tau$ is a finite collection of subsets of Ω, the notion of $\mathcal{G}_\tau$-measurability has an especially simple representation for our purposes. Let $\mathcal{A}_\tau$ denote the subcollection of $\mathcal{G}_\tau$ consisting of all $\mathcal{G}_\tau$-*atoms*, i.e. nonempty $\mathcal{G}_\tau$-measurable sets that do not properly include any other nonempty $\mathcal{G}_\tau$-measurable set. Such atoms are mutually disjoint. A set is $\mathcal{G}_\tau$-measurable if and only if it is a union of $\mathcal{G}_\tau$-atoms. Thus there is a one-to-one correspondence between $\mathcal{G}_\tau$-measurable sets in Ω and sets of $\mathcal{G}_\tau$-atoms, i.e. subsets of $\mathcal{A}_\tau$. A function is $\mathcal{G}_\tau$-measurable if and only if it is constant relative to every $\mathcal{G}_\tau$-atom. Each $\mathcal{G}_\tau$-measurable function can in this way be identified uniquely with a function on $\mathcal{A}_\tau$ rather than on Ω. We can indicate this notationally, when we wish to, by writing $p_{\alpha\tau}$ for $\alpha \in \mathcal{A}_\tau$ to denote the common value that $p_{\omega\tau}$ has for all $\omega \in \alpha$ when p is $\mathcal{G}_\tau$-measurable. (Obviously Ω itself in this setting might be identified with the set of atoms of some finite field of information chosen within a larger, possibly "continuous" probability space by some kind of approximation. We don't go into this matter here.)

Conditional expectation with respect to $\mathcal{G}_\tau$ is denoted by $E^{\mathcal{G}_\tau}$. This can be viewed in the present setting as the linear transformation that takes a random variable such as B_τ and redefines it to have a constant value on each $\mathcal{G}_\tau$-atom $\alpha \in \mathcal{A}_\tau$, that value being, of course, the "weighted average"

$$\left[\sum_{\omega\in\alpha} \pi_\omega B_{\omega\tau}\right] \Big/ \left[\sum_{\omega\in\alpha} \pi_\omega\right].$$

The stochastic dynamical systems for our primal and dual problems are taken again to have the forms (3.1) and (3.2), but with all elements now interpreted as (potentially) random, and with the restriction that

$$u_\tau \text{ is } \mathcal{G}_\tau - \text{measurable}, \tag{4.2}$$

$$v_\tau \text{ is } \mathcal{G}_\tau - \text{measurable}. \tag{4.3}$$

The condition $u_\tau \in U_\tau$ in (3.1) is interpreted to mean that $u_{\omega\tau} \in U_{\omega\tau}$ for all $\omega \in \Omega$, and similarly for $v_\tau \in V_\tau$. Our primal problem of stochastic control is

$$(\mathcal{P}_{\text{sto}}) \quad \begin{aligned} &\text{minimize subject to (3.1) and (4.2) the function } f(u) = \\ &\sum_{\tau=0}^{T} E\{p_\tau{\cdot}u_\tau + \tfrac{1}{2}u_\tau{\cdot}P_\tau u_\tau\} - \sum_{\tau=1}^{T+1} E\{c_\tau{\cdot}x_{\tau-1}\} \\ &+\sum_{\tau=1}^{T} E\{\rho_{V_\tau,Q_\tau}(q_\tau - E^{\mathcal{G}_\tau}\{C_\tau x_{\tau-1}\} - D_\tau u_\tau)\} \\ &\quad + E\{\rho_{V_{T+1},Q_{T+1}}(q_{T+1} - E^{\mathcal{G}_{T+1}}\{C_{T+1}x_T\})\}. \end{aligned}$$

The corresponding dual problem is

$$(\mathcal{Q}_{\text{sto}}) \quad \begin{aligned} &\text{maximize subject to (3.2) and (4.3) the function } g(v) = \\ &\sum_{\tau=1}^{T+1} E\{q_\tau{\cdot}v_\tau - \tfrac{1}{2}v_\tau{\cdot}Q_\tau v_\tau\} - \sum_{\tau=1}^{T} E\{b_\tau{\cdot}y_{\tau+1}\} \\ &-\sum_{\tau=1}^{T} E\{\rho_{U_\tau,P_\tau}(E^{\mathcal{G}_\tau}\{B_\tau^* y_{\tau+1}\} + D_\tau^* v_\tau - p_\tau)\} - E\{\rho_{U_0,P_0}(E^{\mathcal{G}_0}\{B_0^* y_1 - p_0\})\}. \end{aligned}$$

Here ρ_{V_τ,Q_τ} and ρ_{U_τ,P_τ} are "random functions" that depend $\mathcal{G}_\tau$-measurably on $\omega \in \Omega$ by virtue of (4.1). The random variables

$$\xi_\tau := E^{\mathcal{G}_\tau}\{C_\tau x_{\tau-1}\} \text{ and } \eta_\tau := E^{\mathcal{G}_\tau}\{B_\tau^* y_{\tau+1}\}$$

are $\mathcal{G}_\tau$-measurable too, of course, so the arguments to which ρ_{V_τ,Q_τ} and ρ_{U_τ,P_τ} are applied are always $\mathcal{G}_\tau$-measurable. The ρ terms at time τ thus monitor "constraint expressions" based solely on the information available to the decision maker at time τ. Note from the dynamics that $\xi_{\omega\tau}$ depends affinely on $u_{\omega 0}, \ldots, u_{\omega,\tau-1}$, whereas $\eta_{\omega\tau}$ depends affinely on $v_{\omega,\tau+1}, \ldots, v_{\omega,T+1}$.

In order to appreciate the generality of problem $(\mathrm{P}_{\mathrm{sto}})$ it is important, especially for readers accustomed to the traditional approach to stochastic control, to understand the nature of the information structure that is adopted. This structure, which is typical of the literature on stochastic programming, has sometimes been interpreted narrowly as excluding models where the information on which decisions can be based is generated by observations that may be influenced by previous control decisions, cf. the comments of Bertsekas and Shreve [**13**, pp. 10–11]. Such is not actually the case when measurability requirements are referred to a single underlying space, as we shall explain. Thus the specification of the information field $\mathcal{G}_\tau$ as independent of $u_0, u_1, \ldots, u_{\tau-1}$, should *not* be taken to mean, for instance, that in choosing u_τ we are unable to respond to complete or partial observations of the states $x_0, x_1, \ldots, x_{\tau-1}$, inasmuch as those states are generally random variables whose distributions depend on $u_0, u_1, \ldots, u_{\tau-1}$.

The crucial distinction is that of controls u_τ seen directly as functions on the space Ω, rather than controls represented in a feedback mode as functions of past observations and expressible only in a secondary way, through composition, as functions on Ω. The feedback mode of representation, while conceptually very appealing, can be a handicap in our opinion when imposed right from the beginning in the problem formulation. We prefer to proceed at first without it and to recover feedback laws later from optimality conditions, if desired.

Let us imagine, to make this more explicit, that at each time $\tau = 0, 1, \ldots, N$ an observation $z_\tau \in \mathbb{R}^{m_\tau}$ is made before the control decision u_τ is chosen. Of course z_τ is a random variable whose distribution is given by a probability measure μ_τ on $\mathbb{R}^{m_\tau}$, which in general might depend on the controls $u_0, u_1, \ldots, u_{\tau-1}$. Let us suppose that the only information available for the selection of u_τ is the sequence $z_0, z_1, \ldots, z_\tau$. In stochastic control it is common to express this requirement by taking u_τ to be a function of $z_0, z_1, \ldots, z_\tau$, i.e. as a function of a random argument in $\mathbb{R}^{m_0} \times \mathbb{R}^{m_1} \times \cdots \times \mathbb{R}^{m_\tau}$. What we propose instead is to handle $z_0, z_1, \ldots, z_\tau$ as functions defined on the underlying probability space Ω and take u_τ to be a function on Ω that is measurable with respect to the σ-field generated by $z_0, z_1, \ldots, z_\tau$; it is this field that should be identified with $\mathcal{G}_\tau$ in our model. (We have assumed in this paper that Ω is a finite, discrete set, but the idea under consideration applies more generally.) This condition is tantamount to the requirement that u_τ be *representable* by composition of $(z_0, z_1, \dots, z_\tau)$ with some mapping into $\mathbb{R}^{k_\tau}$ from the probability space in $\mathbb{R}^{m_0} \times \mathbb{R}^{m_1} \times \cdots \times \mathbb{R}^{m_\tau}$ induced by these random variables, but it leaves the particular representation open to later investigation.

The advantage to our approach in this setting is that the field $\mathcal{G}_\tau$ may well be independent of $u_0, u_1, \dots, u_{\tau-1}$, even though the distribution of $(z_0, z_1, \dots, z_\tau)$ might not. To this extent we are able to make use of properties of convexity and duality that otherwise could be overlooked.

Before we return to the characterization of optimal controls and trajectories, let us also note that because we allow the dimensionality of the state and control vectors to vary over time, our model also includes classical multistage recourse problems. Suppose that the equations (3.1) have the special form

$$\begin{aligned} x_\tau &= \begin{bmatrix} I \\ 0 \end{bmatrix} x_{\tau-1} + \begin{bmatrix} 0 \\ I \end{bmatrix} u_\tau \qquad \text{for } \tau = 1, \dots, T, \\ x_0 &= u_0, \end{aligned}$$

where the identity matrices I and zero matrices 0 are of the appropriate dimensions. Then

$$x_0 = u_0,\ x_1 = \begin{bmatrix} u_0 \\ u_1 \end{bmatrix},\ x_2 = \begin{bmatrix} u_0 \\ u_1 \\ u_2 \end{bmatrix},\ \text{etc.}$$

Thus x_τ is the "memory" of all decisions up through time τ. Assuming that $\mathcal{G}_{\tau-1} \subset \mathcal{G}_\tau$, we get x_τ, like u_τ, to be $\mathcal{G}_\tau$-measurable. Then in $(\mathcal{P}_{\text{sto}})$ the term

$$q_\tau - E^{\mathcal{G}_\tau}\{C_\tau x_{\tau-1}\} - D_\tau u_\tau$$

represents a general affine expression in $u_0, u_1, \dots, u_\tau$. When ρ_{V_τ, Q_τ} is of the type (2.4), we can rewrite $(\mathcal{P}_{\text{sto}})$ in terms of linear constraints and a quadratic objective involving only the control variables $u_0, u_1, \dots, u_T$. This problem, with its block angular structure, is in the usual format for the multistage stochastic programs with recourse; see [**11**] or [**12**], for example.

Problem $(\mathcal{P}_{\text{sto}})$ revolves around the choice of the random variable $u = (u_0, u_1, \dots, u_T)$, which can be regarded as a function from Ω to $\mathbb{R}^{k_0} \times \cdots \times \mathbb{R}^{k_T}$ and therefore as an element of the finite-dimensional vector space consisting of all such functions. The dimension of this space may be very large indeed just from the size of Ω and possibly T, even if $k_0, \dots, k_T$ are themselves relatively small, as might generally be supposed. We must therefore think of $(\mathcal{P}_{\text{sto}})$ as inherently a "large-scale" problem for which approximate methods of solution will be more appropriate than "exact" ones.

Nevertheless it is well to keep in mind that the representation of u as a function from Ω to $\mathbb{R}^{k_0} \times \cdots \times \mathbb{R}^{k_T}$ tends to exaggerate the dimensionality of ($\mathcal{P}$sto). The constraint that u_τ be $\mathcal{G}_\tau$-measurable means, as already noted, that u_τ can be identified uniquely with a certain function from $\mathcal{A}_\tau$ to $\mathbb{R}^{k_\tau}$. The dimension of the space of all functions from $\mathcal{A}_\tau$ to $\mathbb{R}^{k_\tau}$ is $a_\tau k_\tau$, where

$$a_k = |\mathcal{A}_k| \text{ (the number of atoms in } \mathcal{G}_k).$$

Thus the "true" dimensionality of $(\mathcal{P}_{\text{sto}})$, in the sense of the number of real-valued decision variables, is

$$k^* = a_0 k_0 + a_1 k_1 + \cdots + a_T k_T. \tag{4.5}$$

By the same token, the "true " dimensionality of $(\mathcal{Q}_{\text{sto}})$, where the random variable $v = (v_1, \ldots, v_T, v_{T+1})$ must be optimized, is

$$\ell^* = a_1 \ell_1 + \cdots + a_T \ell_T + a_{T+1} \ell_{T+1}. \tag{4.6}$$

Proposition 4.1. *Let*

$$\mathcal{U} = \{u = (u_0, u_1, \ldots, u_T) |\ u_\tau \textit{ is } \mathcal{G}_\tau\textit{-measurable with } u_\tau \in U_\tau\},$$
$$\mathcal{V} = \{v = (v_1, \ldots, v_T, v_{T+1}) |\ v_\tau \textit{ is } \mathcal{G}_\tau\textit{-measurable with } v_\tau \in V_\tau\},$$

and define $\mathcal{J}(u, v) = E\{J(u, v)\}$, where $J(u, v)$ is the expression in Proposition 3.2 (regarded now as a random variable depending on the choice of the random variables u and v). Then $\mathcal{U}$ and $\mathcal{V}$ are polyhedral convex sets (nonempty), and $\mathcal{J}$ is a quadratic convex-concave function.

Proof. By definition $\mathcal{U}$ is a subset of the space of all functions from Ω to $\mathbb{R}^{k_0} \times \cdots \times \mathbb{R}^{k_T}$ consisting of the functions u such that $u_{\omega\tau} \in U_{\omega\tau}$ for all ω and τ, and $U_{\omega\tau}$ is constant in ω with respect to each $\mathcal{G}_\tau$-atom $\alpha \in \mathcal{A}_\tau$. These conditions can be represented by a finite system of linear equations and inequalities, because Ω is finite and $U_{\omega\tau}$ is by assumption a convex polyhedron for each ω and τ. (Alternatively $\mathcal{U}$ can be viewed as a direct product of polyhedral convex sets $U_{\alpha\tau}$ indexed by $\alpha \in \mathcal{A}_\tau$ and $\tau = 0, 1, \ldots, T$, inasmuch as U_τ is $\mathcal{G}_\tau$-measurable.) Thus $\mathcal{U}$ is a convex polyhedron. Similarly $\mathcal{V}$ is a convex polyhedron. We have by definition

$$\mathcal{J}(u, v) = \sum_{\omega \in \Omega} \pi_\omega J(u_{\omega 0}, u_{\omega 1}, \ldots, u_{\omega T}; v_{\omega 1}, \ldots, v_{\omega T}, v_{\omega, T+1})$$

where the J term for each ω is quadratic convex-concave function and the coefficients π_ω are nonnegative therefore $\mathcal{J}$ is a quadratic convex-concave function. $\square$

Theorem 4.2. *The stochastic optimal control problems* $(\mathcal{P}_{\rm sto})$ *and* $(\mathcal{Q}_{\rm sto})$ *are the primal and dual problems of generalized linear-quadratic programming associated with the* $\mathcal{U},\mathcal{V}$ *and* $\mathcal{J}$ *in Proposition 4.1. In particular, the assertions of Theorem 2.1 are valid for* $(\mathcal{P}_{\rm sto})$ *and* $(\mathcal{Q}_{\rm sto})$.

Proof. We must show that the supremum of $\mathcal{J}(u,v)$ over all $v\in\mathcal{V}$ is the function $f(u)$ in $(\mathcal{P}_{\rm sto})$, and the infimum of $\mathcal{J}(u,v)$ over all $u\in\mathcal{U}$ is $g(u)$ in $(\mathcal{Q}_{\rm sto})$. Starting with $J(u,v)$ in the form of (3.5) (which is obtained by using the right hand expression in (3.3) for $[u,v]$) and taking the expectation, we get by (4.1) that

$$\begin{aligned}\mathcal{J}(u,v) = &\sum_{\tau=0}^{T} E\{p_\tau{\cdot}u_\tau + \tfrac{1}{2}u_\tau{\cdot}P_\tau u_\tau\} - \sum_{\tau=1}^{T+1} E\{c_\tau{\cdot}x_{\tau-1}\}\\ &+\sum_{\tau=1}^{T} E\{[q_\tau - E^{\mathcal{G}_\tau}\{C_\tau x_{\tau-1}\} - D_\tau u_\tau]{\cdot}v_\tau - \tfrac{1}{2}v_\tau{\cdot}Q_\tau v_\tau\}\\ &+E\{[q_{T+1} - E^{\mathcal{G}_{T+1}}\{C_{T+1}x_T\}]{\cdot}v_{T+1} - \tfrac{1}{2}v_{T+1}{\cdot}Q_{T+1}v_{T+1}\}\end{aligned}$$

To maximize this over all $v\in\mathcal{V}$, we must maximize separately in each of the v_τ's subject to v_τ being a $\mathcal{G}_\tau$-measurable function with $v_\tau\in V_\tau$. Denote the random variable $q_\tau - E^{\mathcal{G}_\tau}\{c_\tau x_\tau\} - D_\tau u_\tau$ temporarily by r_τ for $\tau = 1,\dots,T$ and $q_{T+1} - E^{\mathcal{G}_{T+1}}\{C_{T+1}x_T\}$ by r_{T+1}. Then each r_τ is $\mathcal{G}_\tau$-measurable and

$$\begin{aligned}\mathcal{J}(u,v) = &\sum_{\tau=0}^{T} E\{p_\tau u_\tau + \tfrac{1}{2}u_\tau{\cdot}P_\tau u_\tau\} - \sum_{\tau=1}^{T+1} E\{c_\tau{\cdot}x_{\tau-1}\}\\ &\sum_{\tau=0}^{T}\sup_{v_\tau\in\mathcal{V}_\tau} E\{r_\tau{\cdot}v_\tau - \tfrac{1}{2}v_\tau{\cdot}Q_\tau v_\tau\},\end{aligned}$$

where $\mathcal{V}_\tau$ is the set of all $\mathcal{G}_\tau$-measurable v_τ with $v_\tau\in V_\tau$. Since $\mathcal{G}_\tau$-measurable functions can be indexed by $\alpha\in\mathcal{A}_\tau$ in place of $\omega\in\Omega$, as explained above, we can write

$$E\{r_\tau{\cdot}v_\tau - \tfrac{1}{2}v_\tau{\cdot}Q_\tau v_\tau\} = \sum_{\alpha\in\mathcal{A}_\tau}\pi_\alpha[r_{\alpha\tau}{\cdot}v_{\alpha\tau} - \tfrac{1}{2}v_{\alpha\tau}{\cdot}Q_{\alpha\tau}v_{\alpha\tau}],$$

where π_α is the probability of the atom α, i.e.

$$\pi_\alpha = \sum_{\omega\in\alpha}\pi_\omega.$$

The supremum of this expression over all $v_\tau \in \mathcal{V}_\tau$ is

$$\begin{aligned}\sum_{\alpha\in\mathcal{A}_T} \pi_\alpha \sup_{v_{\alpha T}\in V_{\alpha T}} \{r_{\alpha T}{\cdot}v\alpha T - \tfrac{1}{2}v_{\alpha T}{\cdot}Q_{\alpha T}v_{\alpha T}\} \\ = \sum_{\alpha\in\mathcal{A}_T} \pi_\alpha\rho_{V_{\alpha T},Q_{\alpha T}}(r_{\alpha T} = E\{\rho_{V_T,Q_T}(r_\tau)\}.\end{aligned}$$

Thus the supremum of $\mathcal{J}(u,v)$ over $v\in\mathcal{V}$ is

$$\sum_{\tau=0}^{T} E\{p_\tau{\cdot}u_\tau - \tfrac{1}{2}u_\tau{\cdot}P_\tau u_\tau\} - \sum_{\tau=1}^{T} E\{c_\tau{\cdot}x_{\tau-1}\} + \sum_{\tau=1}^{T+1} E\{\rho_{V_\tau,Q_\tau}(r_\tau)\},$$

which from choice of the r_τ's is the objective $f(u)$ in $(\mathcal{P}_{\rm sto})$. The argument that the infimum of $\mathcal{J}(u,v)$ over $u\in\mathcal{U}$ is $g(v)$ in $(\mathcal{Q}_{\rm sto})$ follows the same lines. $\square$

Theorem 4.3. *For the $\mathcal{U}$, $\mathcal{V}$, and $\mathcal{J}$ in Theorem 4.2 one has the following decomposability properties for separate minimization in u or maximization in v. The notation is used that*

$$\begin{aligned}\bar{v}_\tau &= q_\tau - E^{\mathcal{G}_\tau}\{C_\tau\bar{x}_{\tau-1}\} - D_\tau\bar{u}_\tau \text{ for } \tau=1,\ldots,T,\\ \bar{r}_{T+1} &= q_{T+1} - E^{\mathcal{G}_T+1}\{C_{T+1}\bar{x}_T\},\\ \bar{s}_\tau &= E^{\mathcal{G}_\tau}\{B_\tau^*\bar{y}_{\tau+1}\} + D_\tau^*\bar{v}_\tau - p_\tau \text{ for } \tau=1,\ldots,T,\\ \bar{s}_0 &= E^{\mathcal{G}_0}\{B_0^*\bar{y}_1\} - p_0,\end{aligned}$$

where $\bar{u}$ and $\bar{v}$ are elements of $\mathcal{U}$ and $\mathcal{V}$, and $\bar{x}$ and $\bar{y}$ are the corresponding trajectories.

(a) $\tilde{u} \in \operatorname*{argmin}_{u\in\mathcal{U}} \mathcal{J}(u,\bar{v})$ *if and only if*

$$\tilde{u}_{\alpha\tau} \in \partial\rho_{U_{\alpha\tau},P_{\alpha\tau}}(\bar{s}_{\alpha\tau}) = \operatorname*{argmax}_{u_{\alpha\tau}\in U_{\alpha\tau}} \{\bar{s}_{\alpha\tau}{\cdot}u_{\alpha\tau} - \tfrac{1}{2}u_{\alpha\tau}{\cdot}P_{\alpha\tau}u_{\alpha\tau}\}$$

for $\tau=0,1,\ldots,T$ and all $\alpha\in\mathcal{A}_\tau$.

(b) $\tilde{v} \in \operatorname*{argmax}_{v\in\mathcal{V}} \mathcal{J}(\bar{u},v)$ *if and only if*

$$\tilde{v}_{\alpha\tau} \in \partial\rho_{V_{\alpha\tau},P_{\alpha\tau}}(\bar{r}_{\alpha\tau}) = \operatorname*{argmax}_{v_{\alpha\tau}\in V_{\alpha\tau}} \{\bar{r}_{\alpha\tau}{\cdot}v_{\alpha\tau} - \tfrac{1}{2}v_{\alpha\tau}{\cdot}Q_{\alpha\tau}v_{\alpha\tau}\}$$

for $\tau=1,\ldots,T,T+1$ and all $\alpha\in\mathcal{A}_\tau$.

Proof. This combines the argument of Theorem 4.2 with the conjugacy facts noted in the proof of Theorem 3.4. $\square$

Theorem 4.4. *Consider $\mathcal{G}_\tau$-measurable $\bar{u}$, $\bar{v}$, and the corresponding trajectories $\bar{x}$ and $\bar{y}$ determined by (3.1) and (3.2). Define the $\mathcal{G}_\tau$-measurable random variables*

$$\begin{aligned}\bar{p}_\tau &= p_\tau - E^{\mathcal{G}_\tau}\{B_\tau^*\bar{y}_{\tau+1}\} \quad \text{for } \tau=0,1,\ldots,T,\\ \bar{q}_\tau &= q_\tau - E^{\mathcal{G}_\tau}\{C_\tau\bar{x}_{\tau-1}\} \quad \text{for } \tau=1,\ldots,T,T+1.\end{aligned}$$

For each $\tau = 1,\dots,T$ and $\alpha \in \mathcal{A}_\tau$ let $(\overline{\mathcal{P}}_{\alpha\tau})$ and $(\overline{\mathcal{Q}}_{\alpha\tau})$ denote the primal and dual problems of generalized linear-quadratic programming associated with

$$J_{\alpha\tau}(u_{\alpha\tau}, v_{\alpha\tau}) = \overline{p}_{\alpha\tau}{\cdot}u_{\alpha\tau} + \tfrac{1}{2}u_{\alpha\tau}{\cdot}P_{\alpha\tau}u_{\alpha\tau} + \overline{q}_{\alpha\tau}v_{\alpha\tau} - \tfrac{1}{2}v_{\alpha\tau}{\cdot}Q_{\alpha\tau}v_{\alpha\tau} - v_{\alpha\tau}{\cdot}D_{\alpha\tau}u_{\alpha\tau}$$

on $U_{\alpha\tau} \times V_{\alpha\tau}$, namely

$(\overline{\mathcal{P}}_{\alpha\tau})$ *minimize* $\overline{p}_{\alpha\tau}{\cdot}u_{\alpha\tau} + \tfrac{1}{2}u_{\alpha\tau}{\cdot}P_{\alpha\tau}u_{\alpha\tau} + \rho_{V_{\alpha\tau},Q_{\alpha\tau}}(\overline{q}_{\alpha\tau} - D_{\alpha\tau}u_{\alpha\tau})$ *over* $u_{\alpha\tau} \in U_{\alpha\tau}$,

$(\overline{\mathcal{Q}}_{\alpha\tau})$ *maximize* $\overline{q}_{\alpha\tau}{\cdot}v_{\alpha\tau} - \tfrac{1}{2}v_{\alpha\tau}{\cdot}Q_{\alpha\tau}v_{\alpha\tau} - \rho_{U_{\alpha\tau},P_{\alpha\tau}}(D^*_{\alpha\tau}v_{\alpha\tau}\overline{p}_{\alpha\tau})$ *over* $v_{\alpha\tau} \in V_{\alpha\tau}$,

and consider also the problems

$(\overline{\mathcal{P}}_{\alpha 0})$ *minimize* $\overline{p}_{\alpha 0}{\cdot}u_{\alpha 0} + \tfrac{1}{2}u_{\alpha 0}{\cdot}P_{\alpha 0}u_{\alpha 0}$ *over* $u_{\alpha 0} \in U_{\alpha 0}$

for $\alpha \in \mathcal{A}_0$, and

$(\overline{\mathcal{Q}}_{\alpha,T+1})$ *maximize* $\overline{q}_{\alpha,T+1}{\cdot}u_{\alpha,T+1} - \tfrac{1}{2}u_{\alpha,T+1}{\cdot}P_{\alpha,T+1}$ *over* $u_{\alpha,T+1} \in U_{\alpha,T+1}$

for $\alpha \in \mathcal{A}_{T+1}$.

Then a necessary and sufficient condition for $\overline{u}$ and $\overline{v}$ to be optimal solutions to the control problems $(\mathcal{P}_{\rm sto})$ and $(\mathcal{Q}_{\rm sto})$, respectively, is that $\overline{u}_{\alpha\tau}$ should be an optimal solution to the subproblem $(\overline{\mathcal{P}}_{\alpha\tau})$ for every $\alpha \in \mathcal{A}_\tau$ and $\tau = 0,1,\dots,T$, and $\overline{v}_{\alpha\tau}$ should be an optimal solution to the subproblem $(\overline{\mathcal{Q}}_{\alpha\tau})$ for every $\alpha \in \mathcal{A}_\tau$ and $\tau = 1,\dots,T,T+1$.

Proof. The argument imitates the one for Theorem 3.5 but uses the relations in Theorem 4.3. □

References

[1] R.T. Rockafellar, "Generalized linear-quadratic programming and optimal control", SIAM J. Control and Optimization **25**, May, 1987.

[2] 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.

[3] R.T. Rockafellar and R.J-B. Wets, "A Lagrangian finite generation technique for solving linear-quadratic problems in stochastic programming", Math. Programming Studies **28**(1986), 63–93.

[4] R.T. Rockafellar and R.J-B. Wets, "Linear-quadratic programming problems with stochastic penalties: the finite generation algorithm", in V. Arkin, A. Shirayev and R. Wets, eds., *Stochastic Optimization* (Lecture Notes in Control and Informations Sciences), Springer-Verlag, Berlin, 1986, 454–560.

[5] A. King, "An implementation of the finite generation method", in Y. Ermoliev and R.J-B.Wets, eds., *Numerical Techniques for Stochastic Programming*, Springer-Verlag, Heidelberg, 1987, 295–311.

[6] R.W. Cottle, "Symmetric dual quadratic programs", Quart. Appl. Math. **21**(1963), 237-243.

[7] M. Frank and P. Wolfe, "An algorithm for quadratic programming", Naval Res. Log. Quart. **3**(1956), 95-110.

[8] R.T. Rockafellar, *Convex Analysis*, Princeton University Press, Princeton, N.J., 1970.

[9] R. Grinold, "Continuous programming part one: linear objectives", J. Math. Anal. Appl. **37**(1972), 130-141.

[10] D. Bertsekas, *Dynamic Programming and Stochastic Control*, Academic Press, New York, 1976.

[11] M. Eisner and P. Olsen, "Duality for Stochastic Programming Interpreted as L.P. in Lp-Space", SIAM J. Applied Mathematics **28**(1975), 779–792.

[12] R.T. Rockafellar and R. Wets, "The Optimal Recourse Problem in Discrete Time: L^1-Multipliers for Inequality Constraints", SIAM J. on Control and Optimization **16**(1978), 16–36.

[13] D. Bertsekas and S. Shreve, *Stochastic Optimal Control: The Discrete Time Case*, Academic Press, New York, 1978.