MATHEMATICS OF OPERATIONS RESEARCH
Vol. 19, No. 2, May 1994
Printed in U.S.A.

MONOTONE OPTIMAL CONTROL OF PERMUTABLE GSMPs

PAUL GLASSERMAN AND DAVID D. YAO

We consider Markovian GSMPs (generalized semi-Markov processes) in which the rates of events are subject to control. A control is *monotone* if the rate of one event is increasing or decreasing in the number of occurrences of other events. We give general conditions for the existence of monotone optimal controls. The conditions are functional properties for the one-step cost functions and, more importantly, structural properties for the GSMP. The main conditions on costs are submodularity or supermodularity with respect to pairs of events. The key structural condition is *strong permutability*, requiring that the state at any time be determined by the number of events of each type that have occurred, regardless of their order. This permits a reformulation of the original control problem into one based only on event counting processes. This reformulation leads to a unified treatment of a broad class of models and to meaningful generality beyond existing results.

1. Introduction. Without special structure, computation of optimal controls for Markov decision processes is generally infeasible. This motivates investigations into the *form* of optimal policies. Perhaps the simplest form is a threshold policy: below a threshold one action is optimal, above it another is. *Switching curves* characterize a less restrictive class of policies: one action is optimal in states lying below the curve, another is optimal in states lying above it. When optimal actions take values in a continuous set, the natural generalization of a switching-curve policy is a *monotone* optimal control. The monotonicity may be with respect to components of the state or the occurrence of other events. Specific problems for which monotone optimal controls have been identified include those in Bartroli and Stidham (1987), Ghoneim and Stidham (1985), Hajek (1984), Rosberg, Varaiya, and Walrand (1982), Serfozo (1981), and Weber and Stidham (1987), among others. In our setting, as in these, the *rates* of events are controllable. A control is monotone if the rate of each event is increasing or decreasing in the number of occurrences of other events.

We show that monotone optimal controls exist in considerable generality and that conditions for their existence are easily verified. Three rather simple type of hypotheses combine to make optimal controls monotone:

- structural conditions on how the occurrence of events drives the evolution of the state;
- submodularity and supermodularity conditions on the one-step cost function; and,
- inequalities for the one-step cost function at the boundary of the state space.

The main contribution of this paper is the development of the link between monotone optimal controls and the structural conditions. With these conditions in place, we use submodularity much as in Topkis's (1978) theory of ordered optimal solutions, though our setting requires some new results of this type. A step in our analysis extends cost functions beyond boundaries. We use the inequalities mentioned

Received January 25, 1992; revised November 25, 1992.
AMS 1980 subject classification. Primary: 90C39, 90C47, 60K15.
IAOR 1973 subject classification. Main: Control Processes. Cross references: Programming: Dynamic.
OR/MS Index 1978 subject classification. Primary: 121 Dynamic Programming/Semi-Markov.
Key words. Dynamic programming, Markov decision processes, generalized semi-Markov processes, submodularity, antimatroids, monotone controls, switching-curve policies.

0364-765X/94/1902/0449/$01.25

above to ensure that submodularity and supermodularity are preserved by these extensions.

The structural conditions we use are the ones studied in Glasserman and Yao (1992a, b) in an investigation of monotoncity and convexity properties of *uncontrolled* generalized semi-Markov processes (GSMPs). Here, we only consider controlled (countable-state) Markov processes, but the GSMP framework remains valuable: it is precisely the rates of *events*, in the GSMP sense, that are subject to control. Moreover, the GSMP notions of event and event list are central to the conditions we give. Briefly, the structure we use consists of two parts: *noninterruption*, requiring that the occurrence of one event never deactivate another; and *strong permutability*, requiring that the system state be determined by the *number* of occurrences of each event, regardless of their order. These conditions have many ramifications, only some of which are used here; see Glasserman (1991) and Glasserman and Yao (1992a, b).

We show that under strong permutability the vector counting process of events in a Markovian GSMP is itself a Markov process; in fact, this condition is nearly necessary —see Proposition 2.4. Possible values of the vector counting process are called *scores*, and the collection of all such scores is the score space of the GSMP. When strong permutability is combined with noninterruption, it becomes possible to reformulate a Markov decision process for the GSMP on its state space as an equivalent control problem for the counting process on the score space. The modified optimality equation can then be analyzed using the types of techniques employed by Ghoneim and Stidham (1985), Hajek (1984), Serfozo (1981), Weber and Stidham (1987), among others. The transformation to counting processes has several benefits. One is that it allows a unified treatment based on standard conditions of submodularity and supermodularity. Working with a state-space formulation often requires introducing problem-specific functional properties that reduce to submodularity on the score space. More generally, since our controls are event rates, it makes sense to look directly at event processes. In this view, states serve mainly to support the evolution of events. Under strong permutability, the state can still be recovered from the event counting processes so there is no loss.

As shown in Glasserman and Yao (1992b), there is an intimate connection between these structural conditions and certain *antimatroid* properties. This connection is valuable in the present setting as well. One important consequence is that the usual closure under unions of set-system antimatroids translates to closure of the score space under componentwise maximum. In other words, the score space is a join semi-lattice; but it is not, in general, a lattice. This distinction separates our work from earlier treatments of monotone optimal controls, especially Weber and Stidham (1987), where a lattice condition is imposed. This is not merely a mathematical nicety: it is easy to give queueing examples that satisfy only the semi-lattice condition and thus fall outside the scope of existing results; we do so in §§7.3, 7.4 and 10.2. Of course, dropping the assumption of a lattice requires adapting some of Topkis's (1978) results; we do this as well.

We now summarize the organization of the rest of the paper. In §2, we give some background on GSMPs and controlled Markovian GSMPs, then define our structural conditions and describe some of their consequences. In §3, we specify the class of optimal control problems we consider and reformulate them as score-space problems. Section 4 gives the basic monotonicity result for optimal controls when the score space has no boundaries. This rather artificial setting serves as a useful tool in subsequent results. Problems with boundaries are handled in §§5 and 6 using alternative sets of hypotheses. Section 5 uses a condition of Weber and Stidham (1987) that rates be controllable to zero; §6 exploits the antimatroid property to replace this condition with additional restrictions on costs. Section 7 translates score-space results to the original state-space problem and gives examples. In §8 we allow, in addition to rate controls, controls that choose between pairs of events. Connections with the models of Hajek (1984) and Chen, Yang, and Yao (1991) are developed in §9. Finally, in §10, we discuss some weaker structural conditions and give an application to a loss system.

2. A generalized semi-Markov process framework.

2.1. *Markovian GSMPs.* The structure of a GSMP is determined by a generalized semi-Markov *scheme* (GSMS) $\mathcal{G} = (\mathbf{S}, \mathbf{A}, \mathcal{E}, p, \mu)$, where $\mathbf{S}$ is a state space, $\mathbf{A} = \{\alpha_1, \dots, \alpha_m\}$ is a set of *events*, $\mathcal{E}$: $\mathbf{S} \to 2^{\mathbf{A}}$ is an *event-list* mapping, $p = \{p(x'; x, \alpha), x, x' \in \mathbf{S}, \alpha \in \mathcal{E}(x)\}$ is a collection of transition probabilities, and $\mu = \{\mu_\alpha(x), \alpha \in \mathbf{A}, x \in \mathbf{S}\}$ a collection of *speeds*. When the GSMP is in state $x \in \mathbf{S}$, the events in $\mathcal{E}(x) \subseteq \mathbf{A}$ are *active*. For each $\alpha \in \mathcal{E}(x)$, a *clock* runs at rate $\mu_\alpha(x) \geqslant 0$. The initial setting of the clock is random and follows a lifetime distribution characteristic of the associated event α. When the clock for α runs out, the event α is said to occur. The GSMP then moves to state x' from state x with probability $p(x'; x, \alpha)$; each $p(\cdot; x, \alpha)$ is a probability mass function on $\mathbf{S}$. In the new state, clock readings are updated, and new clocks are set for newly active events. See Whitt (1980), and references therein, for more detailed descriptions and constructions in this general setting.

In a Markovian GSMP, all event lifetimes are exponential with mean one, and this allows a straightforward construction. Define an infinitesimal generator Q over $\mathbf{S} \times \mathbf{S}$ by

$$Q(x, y) = \begin{cases} \Sigma_{\alpha \in \mathbf{A}} \mu_\alpha(x) \mathbf{1}\{\alpha \in \mathcal{E}(x)\} p(y; x, \alpha), & \text{if } y \neq x; \\ -\Sigma_{x' \neq x} Q(x, x'), & \text{otherwise.} \end{cases}$$

Requiring

$$\sup_{x \in \mathbf{S}} \max_{\alpha \in \mathbf{A}} \mu_\alpha(x) < \infty$$

makes Q bounded and ensures the existence of a (right-continuous) Markov process $X = \{X_t, t \geqslant 0\}$ on $\mathbf{S}$ with Q as its generator. Generally, we fix the initial state x^0 of X and include it in the specification of the scheme $\mathcal{G}$.

In a *controlled* Markovian GSMP, the control (or policy) is μ, mapping $\mathbf{S} \times \mathbf{A}$ into $\mathbf{R}_+$. For $x \in \mathbf{S}$ and $\alpha \in \mathcal{E}(x)$, $\mu_\alpha(x)$ is the rate at which event α is controlled in state x. If $\alpha \notin \mathcal{E}(x)$, the value of $\mu_\alpha(x)$ has no particular meaning and no effect on the evolution of the process; but, for convenience, we do not restrict the domain of $\mu_{\cdot}(x)$ to $\mathcal{E}(x)$. We sometimes write X^μ for the Markovian GSMP obtained under control μ.

We consider schemes with *deterministic routing*—schemes in which for every $x \in \mathbf{S}$ and every $\alpha \in \mathcal{E}(x)$ there is just one $y \in \mathbf{S}$ with $p(y; x, \alpha) > 0$. In this case, $p(y: x, \alpha) = 1$, and we define $\phi(x, \alpha)$ to be the y for which this holds. Thus, $\phi(x, \alpha)$ is the state to which the process moves from x upon the occurrence of α. If $\phi(x, \alpha) \neq x$ and if there is no other $\beta \in \mathcal{E}(x)$ with $\phi(x, \beta) = \phi(x, \alpha)$, then $Q(x, \phi(x, \alpha)) = \mu_\alpha(x)$. A Markovian GSMP can always be based on a scheme with deterministic routing by, if necessary, introducing new events. But this modification can affect the validity of the structural conditions we discuss next, so the distinction between deterministic and probabilistic routing is important.

2.2. *Strings and scores.* A scheme $\mathscr{G} = (\mathbf{S}, \mathbf{A}, \mathscr{E}, \phi, \mu, x^0)$ determines a *language* $\mathscr{L}$ of *feasible strings* over $\mathbf{A}$. A string is just a finite sequence of elements of $\mathbf{A}$. We call a string $\sigma = \alpha^1 \cdots \alpha^n$ *feasible* if $\alpha^1 \in \mathscr{E}(x^0)$ and if there are states $x^1, \ldots, x^{n-1}$ with $\alpha^i \in \mathscr{E}(x^{i-1})$, $i = 2, \ldots, n$, and $x^i = \phi(x^{i-1}, \alpha^i)$, $i = 1, \ldots, n-1$. By convention, the *empty string* is always feasible. The collection of all feasible strings is $\mathscr{L}$. Extend $\phi(x^0, \cdot)$ recursively to all of $\mathscr{L}$ by letting $\phi(x^0, \cdot)$ applied to the empty string be x^0, and by setting $\phi(x^0, \sigma\alpha) = \phi(\phi(x^0, \sigma), \alpha)$ for $\sigma\alpha \in \mathscr{L}$. Thus, for any $\sigma \in \mathscr{L}$, $\phi(x^0, \sigma)$ is the state reached from x^0 through the sequence of events σ.

To each string σ there corresponds an m-vector ($m = |\mathbf{A}|$), denoted by $[\sigma]$, whose α component $[\sigma]_\alpha$ records the number of instances of α in σ. This vector is the *score* of σ. We use this terminology and notation in analogy with that of Shor et al. (1991). An element of $\mathbf{Z}_+^m$ is a *feasible score* if it is the score of some feasible string. The set

$$\mathscr{N} = \{\mathbf{d} \in \mathbf{Z}_+^m : \exists \sigma \in \mathscr{L}, [\sigma] = \mathbf{d}\}$$

of feasible scores is the *score space* of $\mathscr{G}$.

We can now define key structural properties for schemes. It is convenient to distinguish the case of *state-independent* rates, in which $\mu_\alpha(x) = \mu_\alpha(y)$ for all x, $y \in \mathbf{S}$, for all $\alpha \in \mathbf{A}$.

DEFINITION 2.1. A scheme is

(i) *noninterruptive* if $\{\alpha, \beta\} \subseteq \mathscr{E}(x) \Rightarrow \beta \in \mathscr{E}(\phi(x, \alpha))$;

(ii) *permutable* if for all σ^1, $\sigma^2 \in \mathscr{L}$ we have

$$[\sigma^1] = [\sigma^2] \Rightarrow \mathbf{1}\{\alpha \in \mathscr{E}(x^1)\} = \mathbf{1}\{\alpha \in \mathscr{E}(x^2)\} \quad \text{and} \quad \mu_\alpha(x^1) = \mu_\alpha(x^2)$$

for all $\alpha \in \mathbf{A}$, where $x^i = \phi(x^0, \sigma^i)$. In the state-independent case, this means $\mathscr{E}(x^1) = \mathscr{E}(x^2)$;

(iii) *strongly permutable* if for all σ^1, $\sigma^2 \in \mathscr{L}$, $[\sigma^1] = [\sigma^2] \Rightarrow \phi(x^0, \sigma^1) = \phi(x^0, \sigma^2)$.

In a noninterruptive scheme, the occurrence of one event never deactivates any other event. In a permutable scheme, the event list and clock speeds reached through a (feasible) string of events σ depends on σ only through $[\sigma]$. Under strong permutability, the *state* reached through σ is determined by $[\sigma]$.

Noninterruption is invoked by various authors in various settings; an early use appears in Schassberger (1976). In the terminology of formal languages, (i) makes $\mathscr{L}$ *locally free*; see, e.g., Shor et al. (1991). The combination of properties (i) and (iii) is the *commuting condition* of Glasserman (1991). In the state-independent case, noninterruption and permutability are equivalent to condition (M) (for monotonicity) of Glasserman and Yao (1992a, b). In the general case, (i) and (ii) are implied by condition (SM) (for *speeds* monotonicity) of Glasserman and Yao (1992a).

When conditions (ii) or (iii) apply, we make some notational simplifications. With the initial state x^0 fixed, and with $\mathbf{d} = [\sigma]$ for some $\sigma \in \mathscr{L}$, we sometimes write $\mathscr{E}(\sigma)$ or $\mathscr{E}(\mathbf{d})$ for $\mathscr{E}(\phi(x^0, \sigma))$. Given strong permutability, we sometimes write $\phi(x^0, \mathbf{d})$ for $\phi(x^0, \sigma)$ and $\mu(\mathbf{d})$ for $\mu(\phi(x^0, \mathbf{d}))$.

2.3. *Antimatroid connections.* As discussed in Glasserman and Yao (1992b), the language generated by a noninterruptive, permutable scheme is an *antimatroid with repetition*, in the terminology of Shor et al. (1991). A language $\mathscr{L}$ is an antimatroid (with repetition) if it satisfies the following conditions:

(A1). If $\sigma\alpha \in \mathscr{L}$ then $\sigma \in \mathscr{L}$; i.e., every prefix of a feasible string is feasible (including the empty string).

(A2). If σ, $\sigma\alpha$, and $\sigma\beta$ are in $\mathscr{L}$, with $\alpha \neq \beta$, then $\sigma\alpha\beta \in \mathscr{L}$.

(A3). If σ, $\sigma' \in \mathscr{L}$, with $[\sigma] = [\sigma']$, and $\sigma\alpha \in \mathscr{L}$, then $\sigma'\alpha \in \mathscr{L}$.

Condition (A1) makes the language *left-hereditary* or, equivalently, *prefix-closed*. Conditions (A2) and (A3) are obvious analogs of noninterruption and permutability as applied to schemes. Antimatroids are sometimes called *shelling structures*, *APS-greedoids*, or *interval greedoids without upper bounds*; see Dietrich (1989) for a survey. The qualification "with repetition" allows a symbol to appear more than once in a string. Björner (1985) studies these and related structures.

Shor et al. (1991, Lemma 1.2) provide an alternative characterization of antimatroids through the following *strong exchange* property:

(SE) If $\sigma^1, \sigma^2 \in \mathscr{L}$, then there is a $\sigma^1\sigma \in \mathscr{L}$ with $[\sigma^1\sigma] = [\sigma^1] \vee [\sigma^2]$.

Among left-hereditary languages, the antimatroids are precisely those that satisfy (SE). The language of feasible strings generated by any scheme is automatically left-hereditary, so (SE) distinguishes the languages of noninterruptive, permutable schemes.

The antimatroid property underlies the present analysis through the following result:

PROPOSITION 2.2. *Let $\mathscr{N}$ be the score space of a noninterruptive, permutable scheme.*

(i) *$\mathscr{N}$ is a join semi-lattice; in other words, if $\mathbf{d}^1$ and $\mathbf{d}^2$ are feasible scores, so is $\mathbf{d}^1 \vee \mathbf{d}^2$.*

(ii) *Every score has a unique maximal feasible subscore: for every $\mathbf{d} \in \mathbf{Z}_+^m$, there is a $\hat{\mathbf{d}} \in \mathscr{N}$ satisfying $\hat{\mathbf{d}} \leqslant \mathbf{d}$ and dominating all other feasible scores dominated by $\mathbf{d}$. This maximal feasible subscore is called the* basis *of* $\mathbf{d}$.

PROOF. Part (i) is part (v) of Theorem 2.3 in Glasserman and Yao (1992b), and also an immediate consequence of the strong exchange condition. For part (ii), consider the set of feasible scores dominated by some score $\mathbf{d}$. This set is finite, and it is not empty because the zero vector is a feasible score. Now take the maximum over this set. The resulting score is feasible, by part (i). This is the basis of $\mathbf{d}$. □

Though we usually attach the antimatroid property to $\mathscr{L}$, it applies to the score space as well. To make the connection transparent, think of $\mathbf{d} \in \mathbf{Z}_+^m$ as encoding a multiset containing $\mathbf{d}_\alpha$ copies of α. Then the collection of multisets generated by $\mathscr{N}$ contains the empty set and is *accessible* in the sense of, e.g., Korte and Lovász (1983). Proposition 2.2(i) verifies the remaining antimatroid axiom, closure under unions. With this correspondence, our use of "basis" coincides with its usual meaning for matroids and antimatroids; see the treatment in Dietrich (1989), for example.

In the setting of matroids and (set-system) antimatroids, the *rank* of a set is the cardinality of its basis. The analogous definition in the setting of noninterruptive, permutable schemes (Glasserman and Yao 1992b) sets the rank ρ of a score equal to the sum of the elements of its basis: $\rho(\mathbf{d}) = \Sigma_\alpha \hat{\mathbf{d}}_\alpha$. Let e_α denote the score recording one occurrence of α and no occurrence of any other event. Then the rank function of an antimatroid has a property called *local supermodularity*, namely,

$$\rho(\mathbf{d}) < \rho(\mathbf{d} + e_\alpha), \rho(\mathbf{d} + e_\beta) \Rightarrow \rho(\mathbf{d}) + 1 < \rho(\mathbf{d} + e_\alpha + e_\beta), \qquad \forall \beta \neq \alpha,$$

and this property is easily verified for noninterruptive, permutable schemes in Glasserman and Yao (1992b). Ordinary supermodularity,

$$\rho(\mathbf{d}^1 \vee \mathbf{d}^2) + \rho(\mathbf{d}^1 \wedge \mathbf{d}^2) \geqslant \rho(\mathbf{d}^1) + \rho(\mathbf{d}^2), \tag{1}$$

for all $\mathbf{d}^1, \mathbf{d}^2 \in \mathbf{A}_+^m$, does not hold, in general.

2.4. *Conditions for a lattice.* The failure of ρ to be supermodular mirrors the failure of $\mathcal{N}$ to be closed under $\wedge$; and the gap between $\mathcal{N}$ forming a lattice or merely a semi-lattice separates our work from some of the related literature on monotone optimal control. In particular, Weber and Stidham (1987) (and subsequently Bartroli and Stidham (1987) and Veatch and Wein (1991)) impose "compatibility" conditions that enforce a lattice structure where only a semi-lattice is needed. This unnecessarily narrows the scope of their results, as illustrated in §§7.3, 7.4 and 10.2.

In light of the importance of this distinction, we briefly review the additional structure that makes $\mathcal{N}$ a lattice. The key condition is this: for any $\sigma^1, \sigma^2, \sigma^3 \in \mathcal{L}$,

$$\text{if } [\sigma^3] = [\sigma^1] \wedge [\sigma^2] \quad \text{then } \mathcal{E}(\sigma^1) \cap \mathcal{E}(\sigma^2) \subseteq \mathcal{E}(\sigma^3). \tag{2}$$

When rates are state-independent, this implies permutability. Noninterruption and (2) together are equivalent to condition (CX) (for convexity) of Glasserman and Yao (1992b). The following (and some additional characterizations) are proved in Glasserman and Yao (1992b):

PROPOSITION 2.3. (i) *$\mathcal{N}$ is a lattice if $\mathcal{G}$ is noninterruptive and satisfies* (2).
(ii) *A noninterruptive, permutable scheme satisfies* (2) (*and thus has a lattice score space*) *if and only if its rank function is supermodular.*

Tandem queues, with finite or infinite buffers, fit this framework. See Example 3.7 of Glasserman and Yao (1992b) and §7.2 of this paper.

2.5. *A Markovian implication.* We close this section by connecting the foregoing structural conditions with a stochastic property. This connection opens the way to our analysis of optimal controls. To state it, let $D = \{D_t, t \geqslant 0\}$ with $D_t = (D_{\alpha,t})_{\alpha \in \mathbf{A}}$ be the (right-continuous) event-counting process associated with the Markovian GSMP X. In other words, $D_{\alpha,t}$ is the number of occurrences of the event α in the interval $[0,t]$, and D_0 is the zero vector. The score space $\mathcal{N}$ is the natural state space of D: D never leaves $\mathcal{N}$, and if the $\mu_\alpha(x)$'s are strictly positive, then each element of $\mathcal{N}$ is visited by D with strictly positive probability. We now have

PROPOSITION 2.4. *If $\mathcal{G}$ is permutable then D is a Markov process. If μ is strictly positive and D is Markov, then $\mathcal{G}$ is permutable.*

PROOF. For all $t \geqslant 0$, let $\mathcal{F}_t^D$ be the sigma-algebra generated by $\{D_s, 0 \leqslant s \leqslant t\}$. Clearly, $X_t \in \mathcal{F}_t^D$, since the evolution of events determines the current state. (In fact, X_t depends on the string σ_t of events occurring in $[0,t]$, but not their timing.) Thus, for each $\alpha \in \mathbf{A}$, $D_{\alpha,t}$ has $\mathcal{F}_t^D$-intensity $\mu_\alpha(X_t)\mathbf{1}\{\alpha \in \mathcal{E}(X_t)\}$. If $\mathcal{G}$ is permutable, then this intensity depends on X_t only through D_t, so D is Markov.

Conversely, if D is Markov, then the $\mathcal{F}_t^D$-intensity must be a function of D_t only. In other words, $\mu_\alpha(\phi(x^0,\sigma_t))\mathbf{1}\{\alpha \in \mathcal{E}(\phi(x^0,\sigma_t))\}$ depends on σ_t only through its score D_t. Since $P(\sigma_t = \sigma) > 0$ for every feasible σ, for any $t > 0$, we conclude that $\mu_\alpha(\phi(x^0,\sigma))\mathbf{1}\{\alpha \in \mathcal{E}(\phi(x^0,\sigma))\}$ is a function only of $[\sigma]$ for all $\sigma \in \mathcal{L}$. Under the assumption that μ is strictly positive, this implies permutability. □

Proposition 2.4 allows us to formulate control problems for X as (simpler) control problems for D. As an aside, we note that it provides an interesting parallel to characterizations of when a counting process embeded in a Markov process is Poisson; see Melamed (1982) and Serfozo (1989) and references there.

3. The optimal control problem.

3.1. *State-space formulation.* Consider a controlled Markovian GSMP X^{μ}, as in §2.1. We examine optimal controls under the following objective: when the process is in state x, cost accrues at nonnegative rate $h(x)$; exerting the control $\mu_{\alpha}(x)$ costs $c_{\alpha}(\mu_{\alpha}(x))$. The functions c_{α}, $\alpha \in \mathbf{A}$, are continuous; this ensures that $c_{\alpha}(\mu) + a\mu$ attains a minimum on compact sets, for all $a \in \mathbf{R}$, $\alpha \in \mathbf{A}$. To be precise, we should write $c_{\alpha}(\mu_{\alpha}(x))\mathbf{1}\{\alpha \in \mathscr{E}(x)\}$ for $c_{\alpha}(\mu_{\alpha}(x))$: there is no cost to controlling inactive events. To lighten notation, we omit the indicator function.

Each μ_{α} is constrained to lie in an interval $[\underline{\mu}_{\alpha}, \bar{\mu}_{\alpha}]$ with $0 \leqslant \underline{\mu}_{\alpha} < \infty$. We allow $\underline{\mu}_{\alpha} = \bar{\mu}_{\alpha}$, so uncontrolled events do not require separate treatment. With the rates bounded from above, we can subordinate X^{μ} to a Poisson process in the usual way. Let $\lambda = \sum_{\alpha} \bar{\mu}_{\alpha}$ and let $Y^{\mu} = \{Y_n^{\mu}, n \geqslant 0\}$ be a Markov chain on $\mathbf{S}$ with initial state x^0 and transition matrix

$$P^{\mu}(x, y) = \lambda^{-1} \sum_{\alpha:\, \phi(x, \alpha) = y} \mu_{\alpha}(x)\mathbf{1}\{\alpha \in \mathscr{E}(x)\}.$$

Let N be a rate-λ Poisson process independent of Y^{μ}; then $\{Y_{N_t}^{\mu}, t \geqslant 0\}$ and $\{X_t^{\mu}, t \geqslant 0\}$ are equal in law. Without loss of generality, assume $\lambda = 1$.

The cost functions described above combine with a discount factor r, $0 < r < 1$, to give the finite-horizon objective

$$\min_{\mu} E_x^{\mu}\left[\sum_{i=0}^{n-1} r^i\left\{h(Y_i) + \sum_{\alpha} c_{\alpha}(\mu_{\alpha}(Y_i))\right\}\right], \tag{3}$$

and the infinite-horizon objective

$$\min_{\mu} E_x^{\mu}\left[\sum_{i=0}^{\infty} r^i\left\{h(Y_i) + \sum_{\alpha} c_{\alpha}(\mu_{\alpha}(Y_i))\right\}\right]. \tag{4}$$

Equation (3) defines $V_n(x)$, the minimal n-stage cost-to-go from state x. Since h is nonnegative and the c_{α}'s are continuous, the one-step costs are bounded from below. It follows from general results in, e.g., Whittle (1983, Chapter 26) that V_n is well defined for all n and satisfies the dynamic programming equation $V_{n+1} = TV_n$, where T is the operator on functions from $\mathbf{S}$ to $\mathbf{R}$ defined by

$$Tf(x) = h(x) + \sum_{\alpha} \min_{\underline{\mu}_{\alpha} \leqslant \mu_{\alpha} \leqslant \bar{\mu}_{\alpha}} \{c_{\alpha}(\mu_{\alpha}) + \mu_{\alpha} f(\phi(x, \alpha)) + (\bar{\mu}_{\alpha} - \mu_{\alpha}) f(x)\}. \tag{5}$$

By convention, $\phi(x, \alpha) = x$ if $\alpha \notin \mathscr{E}(x)$. Standard results from dynamic programming (e.g., Whittle 1983) imply that $V(x) = \lim_{n \to \infty} V_n(x)$ equals (4).

3.2. *Score-space formulation.* Suppose, now, that the underlying scheme is strongly permutable. We know from Proposition 2.4 that this makes $\{D_t, t \geqslant 0\}$ a Markov process. For any feasible score $\mathbf{d}$, define $\tilde{h}(\mathbf{d}) = h(\phi(x^0, \mathbf{d}))$. For $i = 1, 2, \ldots$, let D_i be the value of D_t at the ith jump of the dominating Poisson process N. Then corresponding to (3) and (4) we have the objectives

$$\min_{\mu} E_{\mathbf{d}}^{\mu}\left[\sum_{i=0}^{n-1} r^i\left\{\tilde{h}(\mathbf{d} + D_i) + \sum_{\alpha} c_{\alpha}(\mu_{\alpha}(\mathbf{d} + D_i))\right\}\right], \tag{6}$$

and

$$\min_{\mu} E_{\mathbf{d}}^{\mu}\left[\sum_{i=0}^{\infty} r^i \left\{\tilde{h}(\mathbf{d}+D_i) + \sum_{\alpha} c_\alpha(\mu_\alpha(\mathbf{d}+D_i))\right\}\right]. \tag{7}$$

We have written $\mu_\alpha(\mathbf{d}+D_i)$ in place of $\mu_\alpha(\phi(x^0, \mathbf{d}+D_i))$. Under strong permutability, costs and controls can be expressed as functions of scores, rather than states, without changing the problem. Indeed, choosing the score $\mathbf{d}$ in (6) and (7) so that $\phi(x^0, \mathbf{d}) = x$ sets these costs equal to (3) and (4), respectively.

To define the corresponding dynamic programming operator we use an additional definition from Glasserman and Yao (1992b). The *characteristic function* χ of a permutable scheme $\mathscr{G}$ maps $\mathscr{N}$ into $\mathscr{N}$, with α-component

$$\chi_\alpha(\mathbf{d}) = \mathbf{d}_\alpha + \mathbf{1}\{\alpha \in \mathscr{E}(\mathbf{d})\}.$$

Thus, χ_α applied to $\mathbf{d}$ increments $\mathbf{d}_\alpha$ if and only if α is active when the score is $\mathbf{d}$. Various properties of a scheme are reflected in χ. For example, a permutable scheme is noninterruptive if and only if χ is monotone; a noninterruptive, permutable scheme satisfies (2) if and only if χ is supermodular. For these and related results, see Glasserman and Yao (1992b).

Now let $\tilde{V}_n(\mathbf{d})$ be the minimal n-stage cost-to-go from score $\mathbf{d}$; this is (6). Then $\tilde{V}_n$ satisfies the dynamic programming recursion $\tilde{V}_{n+1} = \tilde{T}\tilde{V}_n$, where $\tilde{T}$ is the operator on functions from $\mathscr{N}$ to $\mathbf{R}$ defined by

$$\tilde{T}f(\mathbf{d}) = \tilde{h}(\mathbf{d}) + \sum_{\alpha} \min_{\mu_\alpha} \{c_\alpha(\mu_\alpha) + \mu_\alpha f(\chi_\alpha(\mathbf{d})) + (\bar{\mu}_\alpha - \mu_\alpha) f(\mathbf{d})\}. \tag{8}$$

Inclusion of χ_α is critical: when α is inactive, no control can trigger the occurrence of α. Once again, $\tilde{V}(\mathbf{d}) = \lim_{n\to\infty} \tilde{V}_n(\mathbf{d})$ gives (7).

4. Monotone controls. Our objective is to give conditions under which (6) or (7) admits a *monotone* optimal control; that is, an optimal control μ with the property that

$$\forall \alpha \neq \beta, \qquad \mu_\alpha(\mathbf{d}) \text{ is monotone in } \mathbf{d}_\beta. \tag{9}$$

We say that μ_α is increasing in $\mathbf{d}_\beta$ if $\mu_\alpha(\mathbf{d}) \leqslant \mu_\alpha(\mathbf{d}+e_\beta)$ whenever $\mathbf{d}, \mathbf{d}+e_\beta \in \mathscr{N}$, and decreasing if the inequality is reversed. (Throughout this paper, "increasing" and "decreasing" are used in their weak senses.)

To clarify what role each condition plays, we begin by considering (6) and (7) in the simple but artificial case in which the evolution of D is unrestricted. In other words, there are no boundaries, so D potentially visits any point in $\mathbf{Z}_+^m$. This allows us to omit the function χ in (8) and write simply

$$\tilde{V}_{n+1}(\mathbf{d}) = \tilde{h}(\mathbf{d}) + \sum_{\alpha} \min_{\mu_\alpha} \{c_\alpha(\mu_\alpha) + \mu_\alpha \tilde{V}_n(\mathbf{d}+e_\alpha) + (\bar{\mu}_\alpha - \mu_\alpha)\tilde{V}_n(\mathbf{d})\}. \tag{10}$$

This, in turn, can be rewritten as

$$\tilde{V}_{n+1}(\mathbf{d}) = \tilde{h}(\mathbf{d}) + \sum_{\alpha} \min_{\mu_\alpha} \left\{c_\alpha(\mu_\alpha) + \mu_\alpha\left[\tilde{V}_n(\mathbf{d}+e_\alpha) - \tilde{V}_n(\mathbf{d})\right] + \bar{\mu}_\alpha \tilde{V}_n(\mathbf{d})\right\}. \tag{11}$$

Examining the term in curly braces, we see that the marginal cost of an increase in μ_α

is increasing in $[\tilde{V}_n(\mathbf{d} + e_\alpha) - \tilde{V}_n(\mathbf{d})]$. An optimal control increases as the marginal cost decreases; thus, the optimal μ_α is increasing (decreasing) in $\mathbf{d}_\beta$ if $[\tilde{V}_n(\mathbf{d} + e_\alpha) - \tilde{V}_n(\mathbf{d})]$ is decreasing (increasing) in $\mathbf{d}_\beta$. So, monotonicity of the optimal rate follows from monotonicity of the increments of $\tilde{V}_n$.

Through this observation, submodularity of cost function takes a prominent place, as it has in previous related studies. A function $\tilde{h}$ on the lattice $\mathbf{Z}_+^m$ is supermodular if it has the property featured in (1); it is submodular if the property holds with the inequality reversed. Since $\mathbf{Z}_+^m$ is a finite product of totally ordered sets, it follows from Topkis (1978, Theorem 3.1) that $\tilde{h}$ is submodular if and only if it has *antitone differences*, meaning that for all distinct $\alpha, \beta \in \mathbf{A}$,

$$\tilde{h}(\mathbf{d} + e_\beta + e_\alpha) - \tilde{h}(\mathbf{d} + e_\beta) \leqslant \tilde{h}(\mathbf{d} + e_\alpha) - \tilde{h}(\mathbf{d}),$$

or, equivalently,

$$\tilde{h}(\mathbf{d} + e_\beta + e_\alpha) + \tilde{h}(\mathbf{d}) \leqslant \tilde{h}(\mathbf{d} + e_\alpha) + \tilde{h}(\mathbf{d} + e_\beta). \tag{12}$$

Similarly, the functions with *isotone* differences—those for which (12) holds with the opposite inequality—are precisely the supermodular functions on $\mathbf{Z}_+^m$.

In later sections, cost functions will not necessarily be defined on lattices. We say that a function $\tilde{h}$ on a subset F of $\mathbf{Z}_+^m$ is (α, β)-*submodular* if (12) holds whenever all four points are in S. Define (α, β)-supermodularity analogously.

In order to consider submodularity and supermodularity together, we introduce a further definition. Let $\mathscr{S}_-$ and $\mathscr{S}_+$ be sets consisting of distinct pairs of elements of $\mathbf{A}$ and let $\mathscr{S} = (\mathscr{S}_-, \mathscr{S}_+)$. Every pair of distinct events is in $\mathscr{S}_- \cup \mathscr{S}_+$, but these sets need not be disjoint.

DEFINITION 4.1. A function $\tilde{h}$ on a subset S of $\mathbf{Z}_+^m$ is $\mathscr{S}$-*modular* if it is (α, β)-submodular for all $(\alpha, \beta) \in \mathscr{S}_-$ and (α, β)-supermodular for all $(\alpha, \beta) \in \mathscr{S}_+$. Say that $\tilde{h}$ is $\mathscr{S}$-modular *at* $\mathbf{d}$ if it is $\mathscr{S}$-modular on $\{\mathbf{d}, \mathbf{d} + e_\alpha, \mathbf{d} + e_\beta, \mathbf{d} + e_\alpha + e_\beta\}$ for all distinct α, β.

The path to (9) is now clear. Conditions that make $\tilde{V}_n$, $n = 1, 2, \ldots,$ $\mathscr{S}$-modular make an optimal μ_α increasing in $\mathbf{d}_\beta$ if $(\alpha, \beta) \in \mathscr{S}_-$ and decreasing in $\mathbf{d}_\beta$ if $(\alpha, \beta) \in \mathscr{S}_+$. If (α, β) is in both sets then μ_α is constant in $\mathbf{d}_\beta$. To make $\tilde{V}_n$ $\mathscr{S}$-modular, we assume that $\tilde{h}$ is and show that $\tilde{T}$ preserves $\mathscr{S}$-modularity. For any $n = 1, 2, \ldots,$ and every $\alpha \in \mathbf{A}$, let

$$f_\alpha(\mu_\alpha, \mathbf{d}) = c_\alpha(\mu_\alpha) + \mu_\alpha \tilde{V}_n(\mathbf{d} + e_\alpha) + (\bar{\mu}_\alpha - \mu_\alpha)\tilde{V}_n(\mathbf{d}) \tag{13}$$

$$= c_\alpha(\mu_\alpha) + \mu_\alpha\left[\tilde{V}_n(\mathbf{d} + e_\alpha) - \tilde{V}_n(\mathbf{d})\right] + \bar{\mu}_\alpha \tilde{V}_n(\mathbf{d}). \tag{14}$$

Let $g_\alpha(\mathbf{d}) = \min_{\mu_\alpha} f_\alpha(\mu_\alpha, \mathbf{d})$. Then

$$\tilde{V}_{n+1}(\mathbf{d}) = \tilde{T}\tilde{V}_n(\mathbf{d}) = \tilde{h}(\mathbf{d}) + \sum_\alpha g_\alpha(\mathbf{d}). \tag{15}$$

We show that this transformation preserves $\mathscr{S}$-modularity.

We develop this argument using ideas from Topkis (1978), as do Weber and Stidham (1987); but our use brings some modification, primarily because we consider sub- and supermodularity simultaneously. This leads to more general monotonicity results. Also, since in later sections we do not have a lattice, we use a formulation based, instead, on a subset of a product of totally ordered sets. It appears that Weber and Stidham (1987) impose a lattice condition so that Topkis's analysis applies in its original form; our modification makes this unnecessary.

We state our result in the form most convenient for subsequent application, not necessarily in its fullest generality. To keep the statement generic, we use indices i, j, k rather than α, β. Thus, $\mathscr{S}_-, \mathscr{S}_+$ are now pairs of distinct indices (i, j), $i, j = 1, \ldots, m$, $i \neq j$.

LEMMA 4.2. *Let S be a subset of $\mathbf{Z}_+^m$, f_k a function on $[a, b] \times S$, $a, b \in \mathbf{R}$. Denote a typical element of $[a, b] \times S$ by $(u, d_1, \ldots, d_m)$. Define $g_k(\cdot) = \inf_{u \in [a,b]} f_k(u, \cdot)$. Suppose that*

(i) *for all $u \in [a, b]$, $f_k(u, \cdot)$ is $\mathscr{S}$-modular;*

(ii) *if $(k, j) \in \mathscr{S}_-$, then f_k is submodular in (u, d_j) for all values of its other arguments;*

(iii) *if $(k, j) \in \mathscr{S}_+$, then f_k is supermodular in (u, d_j) for all values of its other arguments.*

Then g_k is $\mathscr{S}$-modular.

PROOF. The proof is similar to that of Topkis (1978, Theorem 4.3). Without loss of generality, assume $j \neq k$. We prove only the case $(i, j) \in \mathscr{S}_-$ and $(k, j) \in \mathscr{S}_-$, the other cases following from essentially the same argument. To see that g_k has the required property, consider points $d, d + e_i, d + e_j$, and $d + e_i + e_j$ in S. Temporarily assume that there exist $u^1, u^2 \in [a, b]$ for which $f_k(u^1, d + e_j) = g_k(d + e_j)$ and $f_k(u^2, d + e_i) = g_k(d + e_i)$; i.e., the infima that define the g_k values are attained. If $u^1 \leqslant u^2$, then

$$\begin{aligned}
&g_k(d) + g_k(d + e_i + e_j) \\
&\qquad \leqslant f_k(u^1, d) + f_k\big(u^2, d + e_i + e_j\big) \\
&\qquad \leqslant f_k\big(u^1, d + e_i\big) + f_k\big(u^1, d + e_j\big) - f_k\big(u^1, d + e_i + e_j\big) + f_k\big(u^2, d + e_i + e_j\big) \\
&\qquad \leqslant f_k\big(u^2, d + e_i\big) + f_k\big(u^1, d + e_j\big) - f_k\big(u^1, d + e_i + e_j\big) + f_k\big(u^1, d + e_i + e_j\big) \\
&\qquad = f_k\big(u^2, d + e_i\big) + f_k\big(u^1, d + e_j\big).
\end{aligned}$$

The first inequality follows from the definition of g_k, the second from condition (i) and the third from condition (ii). We conclude that

$$g_k(d) + g_k\big(d + e_i + e_j\big) \leqslant g_k\big(d + e_i\big) + g_k\big(d + e_j\big); \tag{16}$$

i.e., g_k is (i, j)-submodular. If, instead, $u^1 > u^2$, the argument is the same but starts from $g_k(d) + g_k(d + e_i + e_j) \leqslant f_k(u^2, d) + f_k(u^1, d + e_i + e_j)$. To drop the assumption that the infima over $[a, b]$ are attained, fix an $\epsilon > 0$, let u^1 be any point in $[a, b]$ for which $f_k(u^1, d + e_i) \leqslant g_k(d + e_i) + \epsilon$ and define u^2 analogously. The steps above lead to (16) with 2ϵ added on the right. Since ϵ can be made arbitrarily small (16) follows. □

Comparing this result with Topkis (1978, Theorem 4.3), we see that it weakens the assumption that the domain of f_k is a lattice at the expense of requiring that f_k be defined on a product space $[a, b] \times S$ with $[a, b]$ totally ordered. We now have

THEOREM 4.3. *Suppose $\tilde{h}$ is $\mathscr{S}$-modular. Then for the problem* (10) *without boundaries, under either the finite- or infinite-horizon objective, there exists a monotone* *optimal control* μ. *Specifically, if* $(\alpha, \beta) \in \mathscr{S}_-$ $(\mathscr{S}_+)$ *then* μ_α *is increasing* (*decreasing*) *in* $\mathbf{d}_\beta$.

PROOF. As argued just after (11), to establish that the optimal μ_α increases (decreases) in $\mathbf{d}_\beta$, it suffices to show that every $\tilde{V}_n$, $n \geqslant 0$, is (α, β)-submodular (-supermodular). Since $\tilde{V}_0$ is identically zero, the result is proved once we verify that $\tilde{T}$ preserves $\mathscr{S}$-modularity.

Take as induction hypothesis that $\tilde{V}_n$ is $\mathscr{S}$-modular, and let f_α be as in (13–14). We claim that each f_α satisfies conditions (i)–(iii) of Lemma 4.2, with the index α replacing the index k. Observe in (13) that f_α is a positive linear combination of $\mathscr{S}$-modular functions and is therefore $\mathscr{S}$-modular as well. If $(\alpha, \beta) \in \mathscr{S}_-$, then $[\tilde{V}_n(\mathbf{d} + e_\alpha) - \tilde{V}_n(\mathbf{d})]$ is decreasing in $\mathbf{d}_\beta$, so it follows from (14) that increments of f_α with respect to μ_α are decreasing in $\mathbf{d}_\beta$; they are increasing in $\mathbf{d}_\beta$ if $(\alpha, \beta) \in \mathscr{S}_+$. This verifies properties (ii) and (iii).

It now follows from Lemma 4.2 that each $g_\alpha(\cdot) = \min_{\mu_\alpha} f(\mu_\alpha, \cdot)$ is $\mathscr{S}$-modular. Since $\mathscr{S}$-modularity is preserved under summation, we conclude from (15) that $\tilde{V}_{n+1}$ is $\mathscr{S}$-modular. This completes the induction. □

REMARK. For simplicity, we have assumed that $\mathscr{S}$ contains all pairs of distinct events. However, examination of the proofs of Lemma 4.2 and Theorem 4.3 shows that a weaker condition would suffice. For any event α with $\underline{\mu}_\alpha = \bar{\mu}_\alpha$, the minimization over μ_α i the definition of g_α is superfluous. Consequently, f_α need not satisfy conditions (ii) and (iii) of Lemma 4.2, and for any $\beta, (\alpha, \beta)$ need not be in either $\mathscr{S}_-$ or $\mathscr{S}_+$. Thus, it is enough for $\mathscr{S}$ to contain all pairs of distinct, genuinely controllable events—those for which $\underline{\mu}_\alpha < \bar{\mu}_\alpha$. To simplify the exposition, we continue to assume that $\mathscr{S}$ includes all pairs of distinct events.

5. Penalties at the boundary. We now modify the analysis of the previous section to allow for boundaries on the state space of D. In this section and the one that follows, we present two methods for incorporating boundaries, leading to two sets of hypotheses for monotone controls. Both methods involve extending the function $\tilde{h}$—initially defined only on $\mathscr{N}$—to all of $\mathbf{Z}_+^m$. The extension in this section assigns infinite penalties to infeasible scores; that of §6 projects infeasible scores to their bases.

We begin with a condition of Weber and Stidham (1987) that certain rates be controllable to zero. Call an event *permanent* if it is in the event list of every state, *nonpermanent* otherwise. Our assumption, then, on every $\alpha \in \mathbf{A}$, is

$$\text{if } \alpha \text{ is nonpermanent, then } \underline{\mu}_\alpha = 0. \tag{17}$$

Equivalently, we require that if μ_α cannot be controlled to zero (i.e., if $\underline{\mu}_\alpha > 0$), then α must be active in all states. This condition allows us to replace a problem having boundaries with one not having boundaries by assigning infinite cost to infeasible scores. If $\mathbf{d}$ is feasible and $\alpha \notin \mathscr{E}(\mathbf{d})$, then it follows from the strong exchange property that $\mathbf{d} + e_\alpha$ is infeasible. Thus, if infeasible scores get infinite cost, then any optimal μ must set $\mu_\alpha(\mathbf{d}) = 0$ whenever $\alpha \notin \mathscr{E}(\mathbf{d})$. It follows that an optimal control for the unconstrained problem never allows D to make an infeasible transition, and thus provides an optimal control for the original, constrained problem.

We need to verify that extending $\tilde{h}$ beyond $\mathscr{N}$ in this way preserves $\mathscr{S}$-modularity, the property used in Theorem 4.3. Preservation, we will see, follows from the semi-lattice structure established in Proposition 2.2(i) together with a boundary condition. Let $\bar{h}(\mathbf{d}) = \tilde{h}(\mathbf{d})$ for $\mathbf{d} \in \mathscr{N}$ and set $\bar{h} = \infty$ off $\mathscr{N}$. Say that α *activates* β if there is an $x \in \mathbf{S}$ with $\alpha \in \mathscr{E}(x)$, $\beta \notin \mathscr{E}(x)$, and $\beta \in \mathscr{E}(\phi(x, \alpha))$; i.e., β becomes active just after the occurrence of α in state x. With this terminology, we give the first of several boundary conditions we employ:

(B1) If α activates β, then $\tilde{h}$ is (α, β)-submodular; i.e., $(\alpha, \beta) \in \mathscr{S}_-$.

We call this a boundary condition because if α activates β in score $\mathbf{d}$, then $\mathbf{d}, \mathbf{d} + e_\alpha$, and $\mathbf{d} + e_\alpha + e_\beta$ are in $\mathscr{N}$, but $\mathbf{d} + e_\beta$ is not. We now have

LEMMA 5.1. *Let $\mathscr{N}$ be the score space of a noninterruptive, permutable scheme, and let $\tilde{h}$ be a function on $\mathscr{N}$. Suppose that $\tilde{h}$ is $\mathscr{S}$-modular and that* (B1) *holds. Then $\bar{h}$ is $\mathscr{S}$-modular at* $\mathbf{d}$, *for every* $\mathbf{d} \in \mathscr{N}$. *In particular, if $\tilde{h}$ is submodular then $\bar{h}$ is submodular at* $\mathbf{d}$, *for every* $\mathbf{d} \in \mathscr{N}$.

PROOF. The last statement in the lemma follows from the first part, because if $\tilde{h}$ is submodular then it is $\mathscr{S}$-modular with all distinct pairs of indices in $\mathscr{S}_-$. For the general case, pick any $\mathbf{d} \in \mathscr{N}$ and any $\alpha, \beta \in \mathbf{A}$. We need to verify that $\bar{h}$ has, at $\mathbf{d}$, whatever (α, β)-modularity property $\tilde{h}$ has. If $\mathbf{d} + e_\alpha$ and $\mathbf{d} + e_\beta$ are in $\mathscr{N}$, then so is $\mathbf{d} + e_\alpha + e_\beta$, by the max-closure established for permutable schemes in Proposition 2.2(i). In this case, $\bar{h}$ is $\mathscr{S}$-modular at $\mathbf{d}$ because $\tilde{h}$ is, and because $\bar{h} = \tilde{h}$ on $\mathscr{N}$. Suppose, then, that $\mathbf{d} + e_\beta \notin \mathscr{N}$; i.e., β is not in $\mathscr{E}(\mathbf{d})$, the event list determined by score $\mathbf{d}$. This makes $\bar{h}$ (α, β)-submodular at $\mathbf{d}$, since $\bar{h}(\mathbf{d} + e_\beta) = \infty$; see (12). There are now two cases:

Case (a). Suppose $\mathbf{d} + e_\beta + e_\alpha \in \mathscr{N}$; i.e., there is a $\sigma \in \mathscr{L}$ with $[\sigma] = \mathbf{d} + e_\beta + e_\alpha$. By assumption, $\mathbf{d} \in \mathscr{N}$ so there is a $\sigma^1 \in \mathscr{L}$ with $[\sigma^1] = \mathbf{d}$. By the antimatroid strong exchange property, there is a feasible string $\sigma^1\sigma'$ with score $\mathbf{d} + e_\beta + e_\alpha$; i.e., with $[\sigma'] = e_\beta + e_\alpha$. Since we have assumed $\beta \notin \mathscr{E}(\mathbf{d}) = \mathscr{E}(\sigma^1)$, we must have $\sigma' = \alpha\beta$. In other words, $\mathbf{d} + e_\alpha \in \mathscr{N}$, $\beta \in \mathscr{E}(\mathbf{d} + e_\alpha)$, and α activates β. In this case, $\tilde{h}$ must be (α, β)-submodular, by hypothesis.

Case (b). If $\mathbf{d} + e_\beta + e_\alpha \notin \mathscr{N}$, then $\bar{h}(\mathbf{d} + e_\beta + e_\alpha) = \infty$, so $\bar{h}$ is (α, β)-supermodular as well as (α, β)-submodular, so it automatically matches whichever property $\tilde{h}$ has. □

We can now give our first set of conditions for monotone optimal controls in the presence of boundaries. We formulate the result in terms of the original GSMP control problem.

THEOREM 5.2. *Let X^μ be a controlled Markovian GSMP based on a noninterruptive, strongly permutable scheme. Suppose the rates of nonpermanent events can be controlled to zero, in the sense of* (17). *Suppose also that* (B1) *holds and that $\tilde{h}(\cdot) = h(\phi(x^0, \cdot))$ is $\mathscr{S}$-modular. Then there exists a monotone optimal control μ; μ_α increases or decreases in $\mathbf{d}_\beta$ according as $(\alpha, \beta) \in \mathscr{S}_-$ or $(\alpha, \beta) \in \mathscr{S}_+$.*

PROOF. Once we replace $\tilde{h}$ with $\bar{h}$, we can work with the unconstrained dynamic programming recursion (10), rather than (8). Since $\bar{h}$ is $\mathscr{S}$-modular at $\mathbf{d}$ for every $\mathbf{d} \in \mathscr{N}$, the argument of Theorem 4.3 shows that the restriction of each $\tilde{V}_n$ to $\mathscr{N}$ is $\mathscr{S}$-modular. Monotonicity of μ follows. □

COROLLARY 5.3. *In the setting of Theorem* 5.2, *if $\tilde{h}$ is submodular then there is an optimal control μ such that μ_α increases in $\mathbf{d}_\beta$ for all distinct α, β.*

REMARK. It is important to note that our extension of $\tilde{h}$ beyond $\mathscr{N}$ does not make $\bar{h}$ $\mathscr{S}$-modular on all of $\mathbf{Z}_+^m$; nor does Theorem 5.2 require this. To illustrate, suppose $\tilde{h}$ is submodular on $\mathscr{N}$. Then $\bar{h}$ is submodular except possibly at infeasible scores: if $\mathbf{d}$ is infeasible but $\mathbf{d} + e_\alpha$, $\mathbf{d} + e_\beta \in \mathscr{N}$, then $\bar{h}$ is *super*modular at $\mathbf{d}$. This, however, has no bearing on Theorem 5.2 since we only need to verify that submodularity (more generally, $\mathscr{S}$-modularity) holds at feasible scores.

The extension of $\tilde{h}$ to $\bar{h}$ is indeed submodular on all of $\mathbf{Z}_+^m$ in the presence of the additional structure that makes the score space a lattice. Since other authors have imposed the (stronger-than-necessary) lattice condition, we record the following connection:

PROPOSITION 5.4. *If the scheme $\mathcal{G}$ satisfies* (CX) (*i.e., is noninterruptive and satisfies* (2)), *then $\bar{h}$ is submodular on $\mathbf{Z}_+^m$ if $\tilde{h}$ is submodular on $\mathcal{N}$.*

PROOF. The only case not covered by Lemma 5.1 is $\mathbf{d} \notin \mathcal{N}$, making $\bar{h}(\mathbf{d}) = \infty$. But from Proposition 2.3, we know that $\mathcal{N}$ is a lattice. Thus, either $\mathbf{d} + e_\alpha \notin \mathcal{N}$ or $\mathbf{d} + e_\beta \notin \mathcal{N}$. It follows that $\bar{h}(\mathbf{d} + e_\alpha) + \bar{h}(\mathbf{d} + e_\alpha) = \infty$, so $\bar{h}$ is both submodular and supermodular at $\mathbf{d}$. □

In the Lemma of Weber and Stidham (1987, p. 208), the assumption of submodularity is analogous to requiring $\bar{h}$ to be $\mathcal{S}$-modular on *all* of $\mathbf{Z}_+^m$. In the cases considered by Weber and Stidham, the lattice property follows from a "compatibility" condition. Their examples, as well as those in Bartroli and Stidham (1987) and Veatch and Wein (1992), satisfy the lattice condition (CX) of Glasserman and Yao (1992b); that is, they are noninterruptive and satisfy (2). This is discussed further in §7.

6. Projection to the boundary.

Extending $\tilde{h}$ beyond $\mathcal{N}$ via infinite penalties (and replacing a constrained problem with an unconstrained one) is only effective if rates for inactive events can be controlled to zero; otherwise, all controls yield infinite costs. Thus, to relax (17) we need a different way to handle the boundary. We now extend $\tilde{h}$ beyond $\mathcal{N}$ by projecting infeasible scores to their bases. This motivates

DEFINITION 6.1. For any $\tilde{h}$: $\mathcal{N} \to \mathbf{R}$, the *basis extension* of $\tilde{h}$ is the function $\hat{h}$: $\mathbf{Z}_+^m \to \mathbf{R}$ defined by $\hat{h}(\mathbf{d}) = \tilde{h}(\hat{\mathbf{d}})$, where $\hat{\mathbf{d}}$ is the basis of $\mathbf{d}$.

We use this definition in conjunction with a boundary condition. Say that α activates β in $\mathbf{d}$ if $\alpha \in \mathcal{E}(\mathbf{d})$ and $\beta \in \mathcal{E}(\mathbf{d} + e_\alpha) \setminus \mathcal{E}(\mathbf{d})$. The condition we need is this:

(B2) If α activates β in $\mathbf{d}$ then $\tilde{h}(\mathbf{d} + e_\alpha + e_\beta) \leqslant \tilde{h}(\mathbf{d} + e_\alpha)$.

This condition serves as an alternative to the assumption that rates are controllable to zero. It states that if β is nonpermanent, then the occurrence of β following its activation lowers costs. Paralleling Lemma 5.1, we have

LEMMA 6.2. *Let $\mathcal{N}$ be the score space of a noninterruptive, permutable scheme and let $\tilde{h}$ be a function on $\mathcal{N}$.* (i) *If $\tilde{h}$ is $\mathcal{S}$-modular and satisfies* (B1)–(B2), *then its basis extension $\hat{h}$ is $\mathcal{S}$-modular at* $\mathbf{d}$, *for all* $\mathbf{d} \in \mathcal{N}$. (ii) *If* (B1) *holds and if $\hat{h}$: $\mathbf{Z}_+^m \to \mathbf{R}$ is $\mathcal{S}$-modular at* $\mathbf{d}$ *for all* $\mathbf{d} \in \mathcal{N}$, *then the restriction of $\hat{h}$ to $\mathcal{N}$ is $\mathcal{S}$-modular and satisfies* (B2).

PROOF. We only prove part (i), the argument for part (ii) being similar. Pick any $\mathbf{d} \in \mathcal{N}$ and any $\alpha, \beta \in \mathbf{A}$. Just as in Lemma 5.1, if $\mathbf{d}, \mathbf{d} + e_\alpha, \mathbf{d} + e_\beta$, and $\mathbf{d} + e_\alpha + e_\beta$ are all in $\mathcal{N}$, then $\hat{h}$ inherits $\mathcal{S}$-modularity directly from $\tilde{h}$. Suppose, then that $\beta \notin \mathcal{E}(\mathbf{d})$. As argued in Lemma 5.1, if $\mathbf{d} + e_\alpha + e_\beta$ is feasible, then so is $\mathbf{d} + e_\alpha$. This leaves three cases: of the two points $\mathbf{d} + e_\alpha$ and $\mathbf{d} + e_\alpha + e_\beta$, neither, both, or just $\mathbf{d} + e_\alpha$ may be in $\mathcal{N}$.

Case (a): *Neither*. In this case, all four points $\mathbf{d}, \mathbf{d} + e_\alpha, \mathbf{d} + e_\beta, \mathbf{d} + e_\alpha + e_\beta$, have basis $\mathbf{d}$, so $\hat{h}$ is constant on the four points, and is therefore both (α, β)-submodular and -supermodular.

Case (b): *Just* $\mathbf{d} + e_\alpha$. In this case, $\mathbf{d} + e_\alpha + e_\beta$ has basis $\mathbf{d} + e_\alpha$ (while $\mathbf{d} + e_\beta$ has basis $\mathbf{d}$), so

$$\hat{h}(\mathbf{d}) + \hat{h}(\mathbf{d} + e_\alpha + e_\beta) = \hat{h}(\mathbf{d} + e_\alpha) + \hat{h}(\mathbf{d} + e_\beta),$$

and, once again, $\hat{h}$ is both (α, β)-submodular and -supermodular.

Case (c): *Both*. Since, in this case, α activates β, (B1) requires $(\alpha, \beta) \in \mathscr{S}_-$. Consequently, using (B2) for the inequality, we have

$$\begin{aligned}\hat{h}(\mathbf{d} + e_\alpha + e_\beta) - \hat{h}(\mathbf{d} + e_\alpha) &= \tilde{h}(\mathbf{d} + e_\alpha + e_\beta) - \tilde{h}(\mathbf{d} + e_\alpha)\\ &\leqslant 0\\ &= \tilde{h}(\mathbf{d}) - \tilde{h}(\mathbf{d})\\ &= \hat{h}(\mathbf{d} + e_\beta) - \hat{h}(\mathbf{d});\end{aligned}$$

i.e., $\hat{h}$ is (α, β)-submodular. □

Once we can extend $\tilde{h}$ beyond $\mathscr{N}$ while preserving $\mathscr{S}$-modularity, monotonicity of an optimal control follows:

THEOREM 6.3. *Let X^μ be a controlled Markovian GSMP based on a noninterruptive, strongly permutable scheme. Suppose that $\tilde{h}(\cdot) = h(\phi(x^0, \cdot))$ is $\mathscr{S}$-modular on $\mathscr{N}$ and satisfies* (B1)–(B2). *Then there exists a monotone optimal control μ; μ_α increases or decreases in $\mathbf{d}_\beta$ according as $(\alpha, \beta) \in \mathscr{S}_-$ or $(\alpha, \beta) \in \mathscr{S}_+$.*

PROOF. It is enough to show that $\tilde{V}_n$ is $\mathscr{S}$-modular for all n. Since $\tilde{V}_0$ is identically zero, this is the case if $\tilde{T}$ preserves $\mathscr{S}$-modularity and (B2), where $\tilde{T}$ is the dynamic programming operator in (8). Let $\hat{T}$ be $\tilde{T}$ without boundary constraints; i.e., $\hat{T}$ is the operator in (11) with $\tilde{h}$ replaced by $\hat{h}$. We make the following observations:

(a) for any real-valued function f on $\mathscr{N}$, $\tilde{T}f$ is the restriction of $\hat{T}\hat{f}$ to $\mathscr{N}$, where $\hat{f}$ is the basis extension of f;

(b) if $\hat{f}$ is $\mathscr{S}$-modular at all $\mathbf{d} \in \mathscr{N}$, then so is $\hat{T}\hat{f}$;

(c) if $\hat{T}\hat{f}$ is $\mathscr{S}$-modular at all $\mathbf{d} \in \mathscr{N}$, then $\tilde{T}f$ is $\mathscr{S}$-modular and satisfies (B2).

If (a)–(c) hold and $\tilde{V}_n$ is $\mathscr{S}$-modular and satisfies (B2), then $\hat{V}_n$ is $\mathscr{S}$-modular at all $\mathbf{d} \in \mathscr{N}$ (by Lemma 6.2(i)), and so is $\hat{T}\hat{V}_n$ (by (b)). From (a) and (c), it follows that $\tilde{V}_{n+1}$, the restriction of $\hat{V}_{n+1} = \hat{T}\hat{V}_n$, satisfies $\mathscr{S}$-modularity and (B2). Thus, every $\tilde{V}_n$ is $\mathscr{S}$-modular, if (a)–(c) hold.

Observation (b) is a consequence of Lemma 4.2, and (c) follows from Lemma 6.2(ii) and (a). To verify (a), let $\mathbf{d}$ be a feasible score; then

$$\hat{T}\hat{f}(\mathbf{d}) = \hat{h}(\mathbf{d}) + \sum_\alpha \min_{\mu_\alpha}\left\{c_\alpha(\mu_\alpha) + \mu_\alpha\left[\hat{f}(\mathbf{d} + e_\alpha) - \hat{f}(\mathbf{d})\right] + \bar{\mu}_\alpha \hat{f}(\mathbf{d})\right\}.$$

But if $\mathbf{d}$ is feasible then $\hat{f}(\mathbf{d}) = f(\mathbf{d})$, and $\hat{h}(\mathbf{d}) = \tilde{h}(\mathbf{d})$. Also, the basis of $\mathbf{d} + e_\alpha$ is $\mathbf{d} + e_\alpha$ itself if $\alpha \in \mathscr{E}(\mathbf{d})$ and $\mathbf{d}$ otherwise. Thus, $\hat{f}(\mathbf{d} + e_\alpha) = f(\mathbf{d} + e_\alpha \mathbf{1}\{\alpha \in \mathscr{E}(\mathbf{d})\}) = f(\chi_\alpha(\mathbf{d}))$. With these substitutions, we see from (8) that $\hat{T}\hat{f} = \tilde{T}f$ on feasible scores. □

The following formulation, based on stronger hypotheses, is sometimes convenient.

COROLLARY 6.4. *Let X^μ be a controlled Markovian GSMP based on a noninterruptive, strongly permutable scheme. Suppose that $\tilde{h}$ is submodular on $\mathscr{N}$, and that $\tilde{h}$ is a decreasing function of $\mathbf{d}_\beta$ for every nonpermanent event β. Then there exists a monotone optimal control μ; μ_α increases in $\mathbf{d}_\beta$ for all distinct α, $\beta \in \mathbf{A}$.*

REMARKS. (i) As was the case with $\bar{h}$, the basis extension $\hat{h}$ need not be $\mathscr{S}$-modular on all of $\mathbf{Z}_+^m$. The case in which $\mathbf{d} \notin \mathscr{N}$ is not considered in Lemma 6.2.

(ii) Hajek (1984) establishes the structure of the optimal control for two interacting queues, using steps analogous to (a)–(c) in the proof of Theorem 6.3. In his setting, the state space is $\mathbf{Z}_+^2$, the set of possible queue-length vectors. As a step in his analysis, Hajek extends functions on the state space to all of $\mathbf{Z}^2$ by setting $f(x_1, x_2) = f(x_1^+, x_2^+)$, with $x^+ = \max(x, 0)$. For a simplified version of his model, this coincides with the basis extension when interpreted via scores. To see this, let α_i and δ_i, $i = 1, 2$, denote arrival events at queue i, β_i, $i = 1, 2$, departure events at queue i. (Hajek also allows routing between the two stations.) Then

$$\mathscr{N} = \left\{\mathbf{d} \in \mathbf{Z}_+^6 : \mathbf{d}_{\beta_i} \leqslant \mathbf{d}_{\alpha_i} + \mathbf{d}_{\delta_i}, i = 1, 2\right\}.$$

Projecting an infeasible score $\mathbf{d}$ onto its basis decrements $\mathbf{d}_{\beta_i}$, $i = 1, 2$, sufficiently to meet the constraints in the definition of $\mathscr{N}$. This has the effect of replacing a negative queue length at i with a queue length of zero, which is Hajek's projection. We discuss Hajek's general model and further connections with our results in §§9.1 and 10.1.

To simplify arguments, we have presented the extension $\bar{h}$ of §5 and $\hat{h}$ of this section separately. Theorems 5.2 and 6.3 require that the corresponding conditions of controllability to zero and (B2) hold in all cases. However, these results are easily interwoven to allow some events to satisfy one condition while others satisfy the other.

Let

$$\mathscr{Q} = \left\{\alpha \in \mathbf{A}: \alpha \text{ is nonpermanent and } \underline{\mu}_\alpha = 0\right\}.$$

Extend $\tilde{h}$ to all of $\mathbf{Z}_+^m$ by setting

$$\ddot{h}(\mathbf{d}) = \begin{cases} h(\hat{\mathbf{d}}), & \text{if } \mathbf{d}_\alpha = \hat{\mathbf{d}}_\alpha \quad \text{for all } \alpha \in \mathscr{Q}; \\ \infty, & \text{if } \mathbf{d}_\alpha > \hat{\mathbf{d}}_\alpha \quad \text{for some } \alpha \in \mathscr{Q}. \end{cases}$$

The function $\ddot{h}$ is $\mathscr{S}$-modular at every $\mathbf{d} \in \mathscr{N}$ if all events $\beta \notin \mathscr{Q}$ satisfy (B2). Thus, adapting the proof of Theorem 6.3, we obtain

THEOREM 6.5. *Let X^μ be a controlled Markovian GSMP based on a noninterruptive, strongly permutable scheme. Suppose that $\ddot{h}$ is $\mathscr{S}$-modular on $\mathscr{N}$ and that* (B1) *holds. Suppose further that, for all events β, if $\beta \notin \mathscr{Q}$ then* (B2) *holds. Then there exists a monotone optimal control μ in the sense of Theorem* 6.3.

7. Back to the state space. In this section, through a combination of general observations and specific examples, we examine the state-space implications of the score-space conditions of §§5 and 6. We begin by considering state spaces that are subsets of $\mathbf{Z}^\ell$ for some ℓ.

7.1. *The linear case.* In many permutable queueing systems, the states have a natural encoding as integer vectors and the mapping from scores to states is affine: for $\mathbf{d} \in \mathscr{N}$, $\phi(x^0, \mathbf{d}) = A\mathbf{d} + x^0$. Suppose $\mathbf{S}$ is a subset of $\mathbf{Z}^l$, so A is $l \times m$. There is a direct connection between the columns of A and the event set $\mathbf{A}$. Writing

$$A = \left[\mathbf{a}_1 | \mathbf{a}_2 | \cdots | \mathbf{a}_m\right],$$

we find that, if $\alpha_i \in \mathscr{E}(x)$, then

$$\phi(x, \alpha_i) = x + \mathbf{a}_i. \tag{18}$$

In other words, each event corresponds to a translation in some feasible direction of the state space.

This type of model is the starting point of Weber and Stidham (1987). While many queueing systems fit this framework, not all do. In particular, (18) entails strong permutability because vector addition is commutative. But (18) is more restrictive than strong permutability because it presupposes that the same translation $x \mapsto x + \mathbf{a}_i$ applies throughout the state space. Also, an approach starting from the assumption that $\mathbf{S} \subseteq \mathbf{Z}^l$ and (18) holds, rather than from a GSMP, fails to reveal the structural conditions on the system that make the monotonicity results go through.

When the framework above does apply, through a translation of the state space we may assume that x^0 is the zero vector. In this case, $\tilde{h}(\mathbf{d}) = h(A\mathbf{d})$ for feasible $\mathbf{d}$. Assuming h extends to a function on $\mathbf{R}^l$ that is twice continuously differentiable and writing $\nabla^2 h(x)$ for its matrix of second derivatives evaluated at x, we obtain

$$\frac{\partial \tilde{h}(\mathbf{d})}{\partial \mathbf{d}_{\alpha_i} \partial \mathbf{d}_{\alpha_j}} = \mathbf{a}_i^T \nabla^2 h(A\mathbf{d}) \mathbf{a}_j.$$

If h is convex, then $\nabla^2 h(A\mathbf{d})$ is positive semi-definite; so, if $\mathbf{a}_i$ and $\mathbf{a}_j$ are both directions of increase or both directions of decrease, then the second derivative displayed above is nonnegative and $\tilde{h}$ is (α_i, α_j)-supermodular. If $\mathbf{a}_i$ is a direction of increase but $\mathbf{a}_j$ a direction of decrease (or vice versa), then $\tilde{h}$ is (α_i, α_j)-submodular.

These and related observations further simplify when costs are separable. Suppose that $h(x_1, \dots, x_l) = h_1(x_1) + \cdots + h_l(x_l)$. If x_i is a queue length, then we often have $x_i = \mathbf{d}_{\alpha_i} - \mathbf{d}_{\beta_i}$, for some pair of events α_i, β_i. Suppose no other x_j depends on both α_i and β_i; i.e., only the ith row of A has nonzero entires in both the α_i and β_i columns. Then

$$\frac{\partial \tilde{h}(\mathbf{d})}{\partial \mathbf{d}_{\alpha_i} \partial \mathbf{d}_{\beta_i}} = -h_i''(\mathbf{d}_{\alpha_i} - \mathbf{d}_{\beta_i}),$$

so if h_i is convex, $\tilde{h}$ is (α, β)-submodular, which if h_i is concave, $\tilde{h}$ is (α, β)-supermodular. When $x_i = \mathbf{d}_{\alpha_i} - \mathbf{d}_{\beta_i}$, β_i is typically a nonpermanent event (a service completion) activated by α_i (an arrival from upstream or from an external source). In this case, if (B1) holds, then $\mathscr{S}$-modularity follows from convexity of h_i. Condition (B2) holds under the reasonable additional hypothesis that service completions always lower holding costs. Thus, all the conditions used in §§5 and 6 have natural counterparts when interpreted on the state space. We illustrate them in more detail through specific models.

7.2. *Tandem queues.* This is the example studied in Weber and Stidham (1987). Let l be the number of queues, x_i the queue length at queue i, $i = 1, \dots, l$. An external arrival to the first queue is event β_0; service completion at the ith queue is event β_i, $i = 1, \dots, l$. Take x^0 to be the zero vector and map scores to states via the $l \times (l+1)$ matrix

$$A = \begin{pmatrix} 1 & -1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 1 & -1 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & -1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 1 & -1 \end{pmatrix}.$$

This model is noninterruptive and permutable; indeed, it even satisfies (2) and thus condition (CX) of Glasserman and Yao (1992b). Its score space

$$\mathscr{N} = \left\{ \mathbf{d} \in \mathbf{Z}_+^{l+1} : \mathbf{d}_{\beta_i} - \mathbf{d}_{\beta_{i+1}} \geqslant 0, i = 1, \dots, l \right\}$$

is closed under $\vee$ and $\wedge$.

Let all pairs of distinct events be in $\mathscr{S}_-$, so that $\mathscr{S}$-modularity reduces to submodularity. As noted by Bartroli and Stidham (1987), the functions h for which $h(A\mathbf{d})$ is submodular in $\mathbf{d}$ are precisely the ones Hajek (1985) calls *multimodular*; in light of Hajek's (1985) Proposition 2.2, this characterization can be taken as the definition of multimodularity. Condition (B2) holds if h is increasing in the *partial-sum* order on $\mathbf{Z}^l$; that is, if

$$\sum_{j=1}^{i} x_j \leqslant \sum_{j=1}^{i} y_j, \qquad j = 1, \dots, l \Rightarrow h(x) \leqslant h(y).$$

Under this condition, holding costs decrease as jobs move downstream. With these properties, we summarize some of the consequences of Theorems 5.2 and 6.3 in

PROPOSITION 7.1. *Consider queues in tandem, as above. For both the finite- and infinite-horizon discounted problems, there is an optimal control μ for which μ_{β_i} increases in $\mathbf{d}_{\beta_j}$, for all distinct i, j, under any of the following conditions*:
(i) $\underline{\mu}_{\beta_i} = 0$, $i = 1, \dots, l$, *and h is multimodular*;
(ii) $\overline{\mu}_{\beta_i} = 0$, $i = 1, \dots, l$, *and h is separable with each h_i, $i = 1, \dots, l$, convex*;
(iii) *h is multimodular and increasing in the partial-sum order.*

Part (ii) is the theorem of Weber and Stidham (1987, p. 206), and part (iii) is suggested in a weaker form in the discussion at the end of their paper. The monotonicity with respect to scores asserted in the proposition immediately implies that the optimal service rate at queue i increases in the queue lengths $x_1, \dots, x_i$ and decreases in $x_{i+1}, \dots, x_l$.

Noninterruption and strong permutability continue to hold if the model is modified to allow blocking; *manufacturing*, *communication* and *kanban* blocking, for example, meet the conditions of Glasserman and Yao (1992b, Example 3.7), particularly (CX). Proposition 7.1 holds in that setting as well. Examples of this type appear in Veatch and Wein (1992).

7.3. *Joins and merges.* Consider, now, the system illustrated in Figure 1. There are three queues; α_i denotes arrival to queue i, $i = 1, 2$, and γ denotes service completion at queue 3. Between the first two queues and the third, we consider alternative mechanisms. Under the *join* mechanism, a job at queue 1 is joined to a job at queue 2 when a job is present at both. This "subassembly" then moves to queue 3. In this case, $\beta_1 = \beta_2 = \beta$. Under the *merge* mechanism, jobs move from queue i, $i = 1, 2$, to queue 3 as they are completed, and are served individually at queue 3 in their order of arrival. In this case, β_1 and β_2 are distinct events.

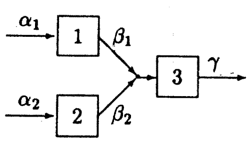

FIGURE 1. When $\beta_1 = \beta_2 = \beta$, jobs are joined before they enter queue 3. When β_1, β_2 are distinct events, the two streams are simply superimposed.

These two mechanisms provide a perfect illustration of the distinction between the lattice and semi-lattice structure. Both systems are noninterruptive and strongly permutable. Under the join mechanism, the score space is

$$\mathscr{N}_{\text{join}} = \left\{ \mathbf{d} \in \mathbf{Z}_+^4 : \mathbf{d}_\beta \leqslant \mathbf{d}_{\alpha_1} \wedge \mathbf{d}_{\alpha_2}; \mathbf{d}_\gamma \leqslant \mathbf{d}_\beta \right\},$$

and this is closed under min as well as max. Building on the compatibility condition of Weber and Stidham (1987), Veatch and Wein (1992) are thus able to describe monotonicity properties of the optimal control for a system much like this one (under the additional assumption that service rates are controllable to zero).

However, when the join becomes a merge, the score space becomes

$$\mathscr{N}_{\text{merge}} = \left\{ \mathbf{d} \in \mathbf{Z}_+^5 : \mathbf{d}_{\beta_i} \leqslant \mathbf{d}_{\alpha_i}, i = 1,2; \mathbf{d}_\gamma \leqslant \mathbf{d}_{\beta_1} + \mathbf{d}_{\beta_2} \right\}.$$

This is closed under max but not min. For example, $(\mathbf{d}_{\alpha_1}, \mathbf{d}_{\beta_1}, \mathbf{d}_{\alpha_2}, \mathbf{d}_{\beta_2}, \mathbf{d}_\gamma) = (1,1,0,0,1)$ and $(0,0,1,1,1)$ are feasible scores, but their min $(0,0,0,0,1)$ is not. In fact, min-closure typically fails whenever there is more than one way to activate some event—in this case γ. See the discussion of minimal elements and unique minimal elements in §§2 and 3 of Glasserman and Yao (1992b).

For simplicity, we consider separable holding costs. Then with the join mechanism, $\tilde{h}(\mathbf{d}) = h_1(\mathbf{d}_{\alpha_1} - \mathbf{d}_\beta) + h_2(\mathbf{d}_{\alpha_2} - \mathbf{d}_\beta) + h_3(\mathbf{d}_\beta - \mathbf{d}_\gamma)$, and with the merge mechanism $\tilde{h}(\mathbf{d}) = h_1(\mathbf{d}_{\alpha_1} - \mathbf{d}_{\beta_1}) + h_2(\mathbf{d}_{\alpha_2} - \mathbf{d}_{\beta_2}) + h_3(\mathbf{d}_{\beta_1} + \mathbf{d}_{\beta_2} - \mathbf{d}_\gamma)$. Let $\mathscr{S}_- = \{(\alpha_1, \beta_1), (\alpha_2, \beta_2), (\alpha_1, \gamma), (\alpha_2, \gamma), (\beta_1, \gamma), (\beta_2, \gamma)\}$ and let $\mathscr{S}_+$ consist of all other pairs of distinct events. From Theorems 5.2 and 6.3, we obtain

PROPOSITION 7.2. *In the model above, suppose* $h(x) = h_1(x_1) + h_2(x_2) + h_3(x_3)$ *with each* h_i *convex. Under the join mechanism, if either*

(i) $\underline{\mu}_\beta = \underline{\mu}_\gamma = 0$, *or*

(ii) h_3 *is increasing and* $h(x_1, x_2, x_3 + 1) \leqslant h(x_1 + 1, x_2 + 1, x_3)$ *for all* x_1, x_2, x_3, *then there is an optimal control* μ *such that* μ_{α_i}, $i = 1,2$, *decrease in* x, μ_β *increases in* x_1 *and* x_2 *and decreases in* x_3, *and* μ_γ *increases in* x. *Under the merge mechanism, if either*

(i′) $\underline{\mu}_{\beta_1} = \underline{\mu}_{\beta_2} = \underline{\mu}_\gamma = 0$, *or*

(ii′) $h(\cdot, x_2, \cdot)$ *and* $h(x_1, \cdot, \cdot)$ *are increasing in the partial-sum order, for all* x_1, x_2, *then there is an optimal control* μ *such that* μ_{α_i}, $i = 1,2$, *decrease in* x, μ_{β_i}, $i = 1,2$, *increase in* x_1 *and* x_2 *and decrease in* x_3, *and* μ_γ *increases in* x.

Since the proofs of Weber and Stidham (1987) and Veatch and Wein (1992) make explicit use of closure under min, it is not clear if their method is extensible to the merge mechanism and similar models, where only the semi-lattice condition applies. Indeed, the weaker conditions developed here appear to be necessary in virtually any model with superimposed arrival streams. Our Proposition 7.2 generalizes in the obvious way to tree networks connecting subsystems like the one in Figure 1. Variations based on Theorem 6.5 are also possible.

7.4. *A re-entrant system.* To further highlight the practical implications of the distinction between the lattice and semi-lattice condition, we consider the control of a simple *re-entrant* system. In a re-entrant *line*, jobs of different classes pass through a series of work centers following deterministic routes; the routes are class-dependent and typically include multiple visits to the same center. Jobs at different stages of their routes waiting for service at a single work center are kept in separate buffers. Thus, each work center provides service to several different buffers. For further discussion of this type of model and its application in semiconductor manufacturing, see Kumar (1993).

For ease of exposition, we consider a single work center with three buffers; the result extends readily to a series of such nodes. Type-1 jobs enter at the first buffer, proceed to the second buffer after their first service completion, then to the third buffer after completion of the second stage of service for a third and final stage. Type-2 jobs enter at the second buffer and subsequently follow the same route as the first class. Type-3 jobs enter at the third buffer and leave upon completing service. At each buffer, service is first come, first served. The service provided to a job depends on the buffer from which it is drawn, but not its class.

Denote the three arrival events by $\alpha_1, \alpha_2, \alpha_3$ and denote completion of service for a buffer-i job by β_i, $i = 1, 2, 3$. The state of the system is (x_1, x_2, x_3) where x_i records the number of jobs in buffer i. This system is noninterruptive and strongly permutable but its score space fails to be closed under min. To see this, suppose the system starts empty and consider the feasible strings $\alpha_1\beta_1\beta_2$ and $\alpha_2\beta_2$. The minimum of their scores corresponds to the string β_2, which is infeasible: no service completion can occur before the first arrival.

As an illustration of the type of conclusion that can still be drawn, consider a separable cost function $h(x) = h_1(x_1) + h_2(x_2) + h_3(x_3)$. Let $\mathscr{S}_-$ consist of all pairs (β_i, β_j), $i \neq j$, and all pairs (α_i, β_j), $i \leqslant j$. Let $\mathscr{S}_+$ contain all other pairs of distinct events. Then we have

PROPOSITION 7.3. *Suppose each h_i, $i = 1, 2, 3$, is convex and suppose all rates are controllable to zero. Then there is an optimal control μ such that μ_{α_i} decreases in x_j, $j \geqslant i$, while μ_{β_i} increases in x_j, $j \leqslant i$ and decreases in x_j for $j > i$.*

This follows from Theorem 5.2. Conclusions based on Theorem 6.3 are also possible.

8. Switch controls.

8.1. *Pure switching.* We now consider systems in which the control μ_α functions as a switch between a pair of events α^1 and α^0. Potential occurrences of either of these events occur at rate $\bar{\mu}_\alpha$; they become potential occurrences of α^1 with probability $\mu_\alpha/\bar{\mu}_\alpha$ and of α^0 with the complementary probability. A potential occurrence is an actual occurrence if the corresponding event is active. With this mechanism, μ_α could control the routing of an arrival stream (in which case α^1, α^0 are arrival events) or it could control the proportion of service given by a single server to two different queues (in which case they are departure events). Every event is now one of a switch-controlled pair $(\alpha^1, \alpha^0), (\beta^1, \beta^0), \ldots$; we still write, e.g., β to refer to a generic event. This model subsumes our previous formulation if we allow $\alpha^0, \beta^0, \ldots$ to be null events that are always active but never change the state.

The state-space dynamic programming recursion for this type of system is

(19)

$$V_{n+1}(x) = h(x) + \sum_{\alpha} \min_{\mu_\alpha} \{c_\alpha(\mu_\alpha) + \mu_\alpha[V_n(\phi(x, \alpha^1)) - V_n(\phi(x, \alpha^0))] + \bar{\mu}_\alpha V_n(\phi(x, \alpha^0))\}.$$

When translated to the score space, the recursion becomes

$$(20)\quad \tilde{V}_{n+1}(\mathbf{d}) = \tilde{h}(\mathbf{d}) + \sum_{\alpha} \min_{\mu_\alpha}\Big\{c_\alpha(\mu_\alpha) + \mu_\alpha\big[\tilde{V}_n(\chi_{\alpha^1}(\mathbf{d})) - \tilde{V}_n(\chi_{\alpha^0}(\mathbf{d}))\big] + \bar{\mu}_\alpha \tilde{V}_n(\chi_{\alpha^0}(\mathbf{d}))\Big\}.$$

Omitting restrictions at the boundary of $\mathscr{N}$, this is

$$(21)\quad \tilde{V}_{n+1}(\mathbf{d}) = \tilde{h}(\mathbf{d}) + \sum_{\alpha} \min_{\mu_\alpha}\Big\{c_\alpha(\mu_\alpha) + \mu_\alpha\big[\tilde{V}_n(\mathbf{d} + e_{\alpha^1}) - \tilde{V}_n(\mathbf{d} + e_{\alpha^0})\big] + \bar{\mu}_\alpha \tilde{V}_n(\mathbf{d} + e_{\alpha^0})\Big\}.$$

Thus, conditions for monotone optimal controls now require that differences $\tilde{V}(\mathbf{d} + e_{\alpha^1}) - \tilde{V}(\mathbf{d} + e_{\alpha^0})$ be monotone in components $\mathbf{d}_\beta$, $\beta \neq \alpha^1, \alpha^0$.

To examine this type of property, let $\mathscr{S} = (\mathscr{S}_-, \mathscr{S}_+)$ be as before, except that now pairs of switch-controlled events (α^1, α^0) need not be included in either $\mathscr{S}_-$ or $\mathscr{S}_+$. Also, for any other event β,

(22) of the pairs $(\alpha^0, \beta), (\alpha^1, \beta)$, one is in $\mathscr{S}_-$, the other in $\mathscr{S}_+$.

In other words, the occurrence of β can increase (decrease) the rate for α^1 only by decreasing (increasing) that for α^0. Consider functions $\tilde{h}\colon \mathbf{Z}_+^m \to \mathbf{R}$ with the following property:

$$(23)\quad \tilde{h}(\mathbf{d} + e_{\alpha^0}) + \tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^1}) \leqslant (\geqslant) \tilde{h}(\mathbf{d} + e_{\alpha^1}) + \tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^0}), \quad \text{if } (\alpha^1, \beta) \in \mathscr{S}_-(\mathscr{S}_+).$$

This reduces to $\mathscr{S}$-modularity when α^0 is a null event. Mimicking the proof of Lemma 4.2, one verifies

LEMMA 8.1. *Let S be a subset of $\mathbf{Z}_+^m$, f_k a function on $[a, b] \times S$, and $g_k(\cdot) = \inf_{u \in [a,b]} f_k(u, \cdot)$. Suppose that f_k satisfies property* (23) *with i^1, i^0, j replacing $\alpha^1, \alpha^0, \beta$. Suppose also that f_k satisfies properties* (ii) *and* (iii) *of Lemma* 4.2. *Then g_k satisfies* (23).

As a consequence, the property in (23) is preserved by value iteration in the absence of boundaries. Thus, Theorem 4.2 continues to hold if $\mathscr{S}$-modularity of $\tilde{h}$ is replaced with (23). In particular, for (21), there is an optimal μ under which μ_α increases in $\mathbf{d}_\beta$ if $(\alpha^1, \beta) \in \mathscr{S}_-$ and decreases in $\mathbf{d}_\beta$ otherwise.

To incorporate boundaries, we use the methods of §§5 and 6. The notions of event permanence and activation are just as in §5. We combine the boundary condition (B2) with the following:

(B3)

(i) if $\alpha^0, \beta \in \mathscr{E}(\mathbf{d})$ and β activates α^1 in $\mathbf{d}$ then

$$\tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^1}) - \tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^0}) \leqslant \tilde{h}(\mathbf{d}) - \tilde{h}(\mathbf{d} + e_{\alpha^0}),$$

(ii) if $\alpha^1, \beta \in \mathscr{E}(\mathbf{d})$ and β activates α^0 in $\mathbf{d}$ then

$$\tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^1}) - \tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^0}) \leqslant \tilde{h}(\mathbf{d} + e_{\alpha^1}) - \tilde{h}(\mathbf{d}).$$

We now have

THEOREM 8.2. *Let X^{μ} be a switch-controlled Markovian GSMP based on a noninterruptive, permutable scheme. Suppose that $\tilde{h}(\cdot) = h(\phi(x^0, \cdot))$ satisfies* (23) *and* (B1) *holds. Then there exists an optimal control μ with μ_α increasing in $\mathbf{d}_\beta$ when $(\alpha^1, \beta) \in \mathscr{S}_-$ and decreasing in $\mathbf{d}_\beta$ when $(\alpha^1, \beta) \in \mathscr{S}_+$ if, for every pair (α^1, α^0) and every other event β any one of the following sets of conditions holds*:

(i) $\underline{\mu}_\alpha = 0$ *and α^0 is permanent*;

(ii) $\overline{\alpha}^1$ *is permanent*;

(iii) (B2) *and* (B3) *hold.*

PROOF. Extend $\tilde{h}$ to $\ddot{h}$ as in Theorem 6.5. If either α^1 or α^0 is permanent and $(\alpha^1, \beta) \in \mathscr{S}_-$, then

$$\ddot{h}(\mathbf{d} + e_\beta + e_{\alpha^1}) + \ddot{h}(\mathbf{d} + e_{\alpha^0}) \leqslant \ddot{h}(\mathbf{d} + e_\beta + e_{\alpha^0}) + \ddot{h}(\mathbf{d} + e_{\alpha^1}),$$

by essentially the argument of Lemma 5.1. If $(\alpha^1, \beta) \in \mathscr{S}_+$, the inequality reverses; thus $\ddot{h}$ shares property (23). For condition (iii), use the argument of Lemma 6.2 to show that (23) extends to $\ddot{h}$. The result now follows much as in Theorem 6.3. □

Condition (B3) simplifies substantially if we strengthen (23) using

LEMMA 8.3. *Suppose that $\tilde{h}$ is $\mathscr{S}$-modular with $\mathscr{S} = (\mathscr{S}_-, \mathscr{S}_+)$ satisfying* (22). *Then $\tilde{h}$ satisfies* (23).

PROOF. Suppose $(\alpha^1, \beta) \in \mathscr{S}_-$, the other case being entirely analogous. Then, if $\alpha^1, \alpha^0, \beta \in \mathscr{E}(\mathbf{d})$,

$$\begin{aligned} &\tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^1}) - \tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^0}) \\ &\quad = \left[\tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^1}) - \tilde{h}(\mathbf{d} + e_\beta)\right] + \left[\tilde{h}(\mathbf{d} + e_\beta) - \tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^0})\right] \\ &\quad \leqslant \left[\tilde{h}(\mathbf{d} + e_{\alpha^1}) - \tilde{h}(\mathbf{d})\right] + \left[\tilde{h}(\mathbf{d}) - \tilde{h}(\mathbf{d} + e_{\alpha^0})\right] \\ &\quad = \tilde{h}(\mathbf{d} + e_{\alpha^1}) - \tilde{h}(\mathbf{d} + e_{\alpha^0}), \end{aligned}$$

where the inequality follows from $\mathscr{S}$-modularity of $\tilde{h}$ and the fact that $(\alpha^0, \beta) \in \mathscr{S}_+$. □

We now obtain the following consequence of Theorem 8.2(iii):

COROLLARY 8.4. *The conclusion of Theorem* 8.2 *holds if conditions* (i)–(iii) *are replaced by the requirement that $\tilde{h}$ be $\mathscr{S}$-modular and satisfy* (B1)–(B2).

PROOF. It suffices to show that (B3) is satisfied. There are several cases; we detail a typical one. Suppose $(\alpha^1, \beta) \in \mathscr{S}_-$, $\alpha^1 \in \mathscr{E}(\mathbf{d} + e_\beta) \setminus \mathscr{E}(\mathbf{d})$ and $\alpha^0 \in \mathscr{E}(\mathbf{d})$. Using first (B2) then the (α^0, β)-supermodularity of $\tilde{h}$, we obtain

$$\begin{aligned} \tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^1}) - \tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^0}) &\leqslant \tilde{h}(\mathbf{d} + e_\beta) - \tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^0}) \\ &\leqslant \tilde{h}(\mathbf{d}) - \tilde{h}(\mathbf{d} + e_\beta + e_{\alpha^0}), \end{aligned}$$

exactly as required by (B3). □

8.2. *Combined switch and rate controls.* In the previous subsection, we controlled the allocation of effort between pairs of events (α^1, α^0), but not the overall level of effort. We now modify the model to include both types of controls. For each pair (α^1, α^0), $p_\alpha \mu_\alpha$ is the rate allocated to α^1, $0 \leqslant p_\alpha \leqslant 1$, and $(1 - p_\alpha)\mu_\alpha$ that allocated to α^0. The state-space dynamic programming recursion becomes

$$V_{n+1}(x) = h(x) + \sum_\alpha \min_{\underline{\mu}_\alpha \leqslant \mu_\alpha \leqslant \bar{\mu}_\alpha} \min_{0 \leqslant p_\alpha \leqslant 1} \{c_\alpha(\mu_\alpha) + \mu_\alpha[p_\alpha V_n(\phi(x, \alpha^1)) + (1 - p_\alpha)V_n(\phi(x, \alpha^0))] + (\bar{\mu}_\alpha - \mu_\alpha)V_n(x)\}. \tag{24}$$

Translating this to the score space, rearranging terms and omitting the boundary restrictions, we obtain

$$\tilde{V}_{n+1}(\mathbf{d}) = \tilde{h}(\mathbf{d}) + \sum_\alpha \min_{\mu_\alpha, p_\alpha} \{c_\alpha(\mu_\alpha) + \mu_\alpha\{p_\alpha[\tilde{V}_n(\mathbf{d} + e_{\alpha^1}) - \tilde{V}_n(\mathbf{d})] + (1 - p_\alpha)[\tilde{V}_n(\mathbf{d} + e_{\alpha^0}) - \tilde{V}_n(\mathbf{d})]\} + \bar{\mu}_\alpha \tilde{V}_n(_h)\}. \tag{25}$$

For simplicity, throughout this section we assume $\underline{\mu}_\alpha = 0$ whenever either of α^1, α^0 is nonpermanent; this is essentially the controllability-to-zero condition (17). We continue to assume that $\mathscr{S} = (\mathscr{S}_-, \mathscr{S}_+)$ satisfies (22); thus, if $\tilde{V}_n$ is $\mathscr{S}$-modular and $(\alpha^1, \beta) \in \mathscr{S}_-$, then

$$\Delta_{\alpha^1}\tilde{V}_n(\mathbf{d}) \triangleq \tilde{V}_n(\mathbf{d} + e_{\alpha^1}) - \tilde{V}_n(\mathbf{d})$$

decreases in $\mathbf{d}_\beta$ and

$$\Delta_{\alpha^0}\tilde{V}_n(\mathbf{d}) \triangleq \tilde{V}_n(\mathbf{d} + e_{\alpha^0}) - \tilde{V}_n(\mathbf{d})$$

increases. Consequently, the combined coefficient on p_α in (25) decreases in $\mathbf{d}_\beta$ and the optimal p_α increases. This, and the behavior of the optimal μ_α, are summarized in

THEOREM 8.5. *For the optimal control problem* (24), *if* $\tilde{h}$ *is* $\mathscr{S}$*-modular with* $\mathscr{S}$ *satisfying* (B1) *and* (22), *then there is an optimal* (μ, p) *such that*

(i) p_α *increases in* $\mathbf{d}_\beta$ *if* $(\alpha^1, \beta) \in \mathscr{S}_-$ *and decreases in* $\mathbf{d}_\beta$ *if* $(\alpha^1, \beta) \in \mathscr{S}_+$;

(ii) *if* $(\alpha^1, \beta) \in \mathscr{S}_-$, *then* μ_α *is first increasing then decreasing in* $\mathbf{d}_\beta$, *while if* $(\alpha^1, \beta) \in \mathscr{S}_+$, *then* μ_α *is first decreasing then increasing in* $\mathbf{d}_\beta$.

PROOF. For part (i), let $\bar{V}_n$ be the extension of $\tilde{V}_n$ obtained when $\tilde{h}$ is extended to $\bar{h}$; $\bar{V}_n$ is infinite on infeasible scores. By the argument of Lemma 5.1, together with the fact that $\mathscr{S}$-modularity is preserved by convex combinations, we conclude that $\bar{V}_n$ is $\mathscr{S}$-modular, so the argument of Theorem 5.2 concludes the proof of part (i). For part (ii), there are three regimes: $\Delta_{\alpha^0}\tilde{V}_n(\mathbf{d}) \leqslant \min(0, \Delta_{\alpha^1}\tilde{V}_n(\mathbf{d}))$; $0 \leqslant \min(\Delta_{\alpha^1}\tilde{V}_n(\mathbf{d}), \Delta_{\alpha^0}\tilde{V}_n(\mathbf{d}))$; and $\Delta_{\alpha^1}\tilde{V}_n(\mathbf{d}) \leqslant \min(0, \Delta_{\alpha^0}\tilde{V}_n(\mathbf{d}))$. In the first, μ_α decreases in $\mathbf{d}_\beta$; in the second, $\mu_\alpha = 0$, and in the third, μ_α increases in $\mathbf{d}_\beta$. By $\mathscr{S}$-modularity, the increments $\Delta_{\alpha^1}\tilde{V}_n$ and $\Delta_{\alpha^0}\tilde{V}_n$ are monotone in opposite directions as functions of $\mathbf{d}_\beta$. Thus, they cross at most once. The order in which they may cross is determined by which of $\mathscr{S}_-$ and $\mathscr{S}_+$ contains (α^1, β) and gives the change of direction indicated in (ii). □

REMARKS. (i) Since we have assigned no cost to exercising the p_α's, an optimal allocation need only use the extreme values 0 and 1. As a function of $\mathbf{d}_\beta$ (with all other $\mathbf{d}_\gamma$'s held fixed), an optimal p_α switches from 0 to 1 where $\Delta_{\alpha^1}\tilde{V}_n(\mathbf{d})$ crosses

$\Delta_{\alpha^0}\tilde{V}_n(\mathbf{d})$. If these increments never cross, then p_α is either identically 0 or identically 1. If we drop the cost c_α, then μ_α may be restricted to its extremes, 0 and $\bar{\mu}_\alpha$, and the following rule is optimal:

$$\Delta_{\alpha^0}\tilde{V}_n(\mathbf{d}) \leqslant \min\left(0, \Delta_{\alpha^1}\tilde{V}_n(\mathbf{d})\right) \Rightarrow \mu_\alpha = \bar{\mu}_\alpha, \, p_\alpha = 0;$$

$$0 \leqslant \min\left(\Delta_{\alpha^1}\tilde{V}_n(\mathbf{d}), \Delta_{\alpha^0}\tilde{V}_n(\mathbf{d})\right) \Rightarrow \mu_\alpha = 0;$$

$$\Delta_{\alpha^1}\tilde{V}_n(\mathbf{d}) \leqslant \min\left(0, \Delta_{\alpha^0}\tilde{V}_n(\mathbf{d})\right) \Rightarrow \mu_\alpha = \bar{\mu}_\alpha, \, p_\alpha = 1.$$

Ties can be broken arbitrarily. Depending on the values of $\mathbf{d}_\gamma$'s, $\gamma \neq \beta$, all three regimes or any subset may be present. However, if $(\alpha, \beta) \in \mathscr{S}_-$ they can only appear in the order listed, as $\mathbf{d}_\beta$ increases, because of the monotonicity of $\Delta_{\alpha^1}\tilde{V}_n$ and $\Delta_{\alpha^0}\tilde{V}_n$. If $(\alpha, \beta) \in \mathscr{S}_+$, the cases listed can only appear in the reverse order. Thus, in either instance there are potentially two switch points. If any one of the regimes never prevails, there is at most one switch point and μ_α is monotone.

(ii) Since the theorem requires that (B1) hold for both α^1 and α^0, and since (22) requires that either (α^1, β) or (α^0, β) belong to $\mathscr{S}_+$, we cannot allow an event β to activate both α^1 and α^0. Different boundary conditions could accomodate that case.

9. Examples with switch controls.

9.1. *Two interacting service stations.* We now apply the foregoing results to a pair of queues sharing a server and an arrival stream; see Figure 2. This is a simplification of Hajek's (1984) model. We return to the general model in §10.

The events are illustrated in Figure 2. (We use superscripts in (β^0, β^1) and (α^0, α^1) to indicate a switch-controlled pair, subscripts in δ_1, δ_2 to indicate unpaired events.) This system is noninterruptive and strongly permutable. For the most natural monotonicity results, and to be consistent with (B1), we specify that $(\alpha^1, \beta^1), (\delta_1, \beta^1), (\alpha^1, \delta_2), (\delta_2, \beta^0)$ are in $\mathscr{S}_-$ and $(\alpha^1, \beta^0), (\delta_1, \beta^0), (\alpha^1, \delta_1), (\alpha^0, \delta_2)$ are in $\mathscr{S}_+$. Under (22), this completely determines $\mathscr{S}$, except for the pair (δ_1, δ_2). Let us initially take $(\delta_1, \delta_2) \in \mathscr{S}_+$.

Letting x_i be the number of jobs at queue i, $i = 1, 2$, the state space becomes $\mathbf{Z}_+^2$. The mapping from scores to states sets $x_1 = \mathbf{d}_{\delta_1} + \mathbf{d}_{\alpha^1} - \mathbf{d}_{\beta^1}$ and $x_2 = \mathbf{d}_{\delta_2} + \mathbf{d}_{\alpha^0} - \mathbf{d}_{\beta^0}$. As in Hajek (1984), consider the functions h on $\mathbf{Z}_+^2$ satisfying the following properties:

(a) h is increasing;

(b) h is convex in each argument and supermodular;

(c1) $h(x_1, x_2 + 1) - h(x_1 + 1, x_2)$ is decreasing in x_1;

(c2) $h(x_1, x_2 + 1) - h(x_1 + 1, x_2)$ is increasing in x_2.

We claim that if h satisfies (a)–(c2), then $\tilde{h}$ meets the conditions of Theorem 8.2(iii): it satisfies (23), (B2) and (B3). There are many cases; we examine some representative ones. Condition (B2) applied to (α^1, β^1) or (δ_1, β^1) requires $h(0, x_2) \leqslant$

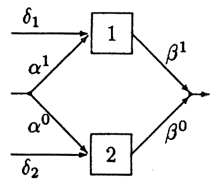

FIGURE 2. The model of §9.1.

$h(1, x_2)$ and applied to (α^0, β^0) or (δ_2, β^0) requires $h(x_1, 0) \leqslant h(x_1, 1)$. These requirements are satisfied by an increasing h. Condition (23) applied to (δ_1, β^1) (with $\delta_1^1 = \delta_1$ and δ_1^0 a null event) requires

$$2h(x_1, x_2) \leqslant h(x_1 + 1, x_2) + h(x_1 - 1, x_2),$$

which is convexity in x_1. The same condition applied to (δ_2, β^0) is satisfied when h is convex in x_2. Supermodularity of h implies

$$h(x_1 + 1, x_2 + 1) + h(x_1, x_2) \geqslant h(x_1 + 1, x_2) + h(x_1, x_2 + 1),$$

which is (23) applied to $(\delta_1, \delta_2) \in \mathscr{S}_+$. Condition (B3)(i) applied to (α^1, β^1) or (δ_1, β^1) requires

$$h(0, x_2) - h(1, x_2 - 1) \leqslant h(0, x_2) - h(0, x_2 - 1),$$

which follows from (a) and (c1). Also, (a) and (c2) imply that (B3)(ii) holds for both (α^0, β^0) and (δ_2, β^0). Though (23) and (B2)–(B3) hold, condition (a)–(c2) do *not*, in general, make $\tilde{h}$ $\mathscr{S}$-modular.

Interpreting the score monotonicity provided by Theorem 8.2(iii) in terms of states, we draw the following conclusions about an optimal μ: μ_α and μ_{δ_1} increase in x_2 and decrease in x_1, while μ_β and μ_{δ_2} increase in x_1 and decrease in x_2. When there is no cost to control (every c_α is identically zero) and when $\mu_{\delta_i} = \bar{\mu}_{\delta_i}$, $i = 1, 2$, these are Hajek's (1984) results applied to our simplified model. Our conditions on $\tilde{h}$ are weaker than conditions (a)–(c2) used in Hajek (1984).

If we move (δ_1, δ_2) from $\mathscr{S}_+$ to $\mathscr{S}_-$ we obtain results outside the scope of Hajek's analysis. In this setting, it is natural to take $\underline{\mu}_{\delta_i} = \bar{\mu}_{\delta_i}$, $i = 1, 2$; otherwise, arrivals to one queue accelerate arrivals to the other. Suppose $\tilde{h}$ satisfies (a)–(c2) but with supermodularity in (b) replaced by submodularity. Then $\tilde{h}$ is $\mathscr{S}$-modular and satisfies (B2), so Corollary 8.4 applies. The monotonicity conclusions drawn above for μ_α and μ_β continue to hold. An example of an increasing, submodular cost function, convex in each coordinate, is $h(x_1, x_2) = \max(x_1, x_2)$.

Hajek's general model includes two features not present in Figure 2: uncontrolled departures from each queue, and controlled routing between the two queues. The first feature poses no problem. If uncontrolled departures from queue i occur at rate ν_i, $i = 1, 2$, then we increase $\bar{\mu}_\beta$ by $\nu_1 + \nu_2$ and constrain μ_β so that $\mu_\beta \geqslant \nu_1$ and $(\bar{\mu}_\beta - \mu_\beta) \geqslant \nu_2$. This incorporates minimal departure rates at the two stations without introducing additional events.

To include routing between the stations, it appears to be necessary to introduce new events—events that violate noninterruption. Let γ_1 and γ_2 denote, respectively, transition from queue 1 to queue 2 and from queue 2 to queue 1. Then γ_1 and β^1 interrupt each other when $x_1 = 1$, as do γ_2 and β^0 when $x_2 = 1$. The score space for this model is

$$(26) \quad \mathscr{N} = \left\{\mathbf{d} \in \mathbf{Z}_+^8 : \mathbf{d}_{\beta^1} + \mathbf{d}_{\gamma_1} \leqslant \mathbf{d}_{\delta_1} + \mathbf{d}_{\alpha^1} + \mathbf{d}_{\gamma_2}, \mathbf{d}_{\beta^0} + \mathbf{d}_{\gamma_2} \leqslant \mathbf{d}_{\delta_2} + \mathbf{d}_{\alpha^0} + \mathbf{d}_{\gamma_1}\right\}.$$

This score space is not max-closed and does not satisfy the strong exchange property. It does, however, satisfy a weaker condition discussed in §10.

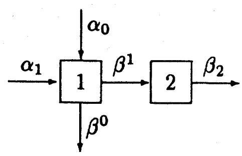

FIGURE 3. The model of §9.2.

9.2. *A scheduling model.* Chen, Yang, and Yao (1991) study the model in Figure 3. There are two classes of arrivals, represented by events α_0 and α_1. Corresponding to these two classes of jobs there are two service-completion events, β^0 and β^1. Upon moving from station 1 to station 2, a class 1 job becomes a class 2 job. Service completions at station 2 are represented by β_2. A system state is a vector $(x_0, x_1, x_2) \in \mathbf{Z}_+^3$, with each x_i recording the number of class i jobs present. This system is easily seen to be noninterruptive and strongly permutable.

The server at station 1 provides an overall service rate μ_β to be allocated between classes 0 and 1. Let this service rate and also μ_{β_2} be controllable to zero. Depending on whether or not μ_{α_0} and μ_{α_1} are controllable, we obtain different conditions and results. Let us take these arrival rates to be fixed, as in Chen, Yang, and Yao (1991, §4). Following the remark at the end of §4, we may restrict modularity conditions on $\tilde{h}$ to pairs of controlled events β^0, β^1, and β_2. Since β^1 activates β_2, we require (β^1, β_2)-submodularity. This entails (β^0, β_2)-supermodularity, via (22) restricted to $\{\beta^0, \beta^1, \beta_2\}$. A function $\tilde{h}(\mathbf{d}) = h(\mathbf{d}_{\alpha_0} - \mathbf{d}_{\beta^0}, \mathbf{d}_{\alpha_1} - \mathbf{d}_{\beta^1}, \mathbf{d}_{\beta^1} - \mathbf{d}_{\beta_2})$ satisfies these conditions if h is convex in x_2, supermodular in (x_0, x_2) and submodular in (x_1, x_2). These are the conditions of Remark (iii) in Chen, Yang, and Yao (1991, §4). They put us in the setting of Theorem 8.5. From part (i) of that result it follows that the optimal fraction of service allocated to class 1 increases in x_1 and decreases in x_0 and x_2, as noted by Chen, Yang, and Yao (1991). Moreover, since no cost is imposed on the allocation, the entire service effort at station 1 can be devoted to a single class in each state without loss of optimality. Thus, in each state, either class 0 or class 1 is given strict priority over the other.

In the case of (increasing) linear holding costs, Chen, Yang, and Yao (1991, Theorem 4.1) show that $\Delta_{\beta^0}\tilde{V}_n(\mathbf{d}) \leqslant 0$ for all feasible $\mathbf{d}$. If there is no cost to providing service, it follows (see Remark (i) of §8.2, above) that μ_β, the optimal overall service effort at station 1 is always $\bar{\mu}_\beta$ if there is a job present at station 1. However, if x_2 is sufficiently large, even if $x_0 = 0$ all effort at station 1 may be allocated to class 0. In other words, for sufficiently large x_2, the optimal policy idles the server at station 1 if no class 0 jobs are present. These conclusions are summarized in Theorem 4.4 of Chen, Yang, and Yao (1991).

10. Monotone control under mutual interruption. Given permutability, max-closure of the score space is equivalent to noninterruption; hence, relaxing the requirement of noninterruption while preserving monotonicity of optimal controls means relaxing max-closure. Relaxing noninterruption is useful in modeling; for example, in §9.1 we noted that introducing switching between stations in the general model of Hajek (1984) violates noninterruption.

In this section, we extend some of our earlier results to systems satisfying *mutual interruption*. We say that α and β interrupt each other in state x if $\alpha, \beta \in \mathscr{E}(x)$ while $\beta \notin \mathscr{E}(\phi(x, \alpha))$ and $\alpha \notin \mathscr{E}(\phi(x, \beta))$. A scheme satisfies mutual interruption if, for every pair of distinct events α, β and every state x, either α and β interrupt each other in state x or else neither interrupts the other. Thus, a noninterruptive scheme is a special case of a mutually interruptive one. If we enhance the model of Figure 2 by including a switching event γ_1 from queue 1 to queue 2, then β^1 and γ_1 interrupt each other when there is just one job at queue 1. Similarly, adding a switching event γ_2 from queue 2 to queue 1 violates noninterruption but satisfies mutual interruption.

10.1. *Weak exchange and its consequences.* Recall from §2.3 that, taken together, permutability and noninterruption are equivalent to the antimatroid strong exchange condition (SE). To accomodate mutual interruption, we relax (SE) to get the following *weak exchange* property:

(WE) If $\sigma^1, \sigma^2 \in \mathscr{L}$ and $[\sigma^1] \leqslant [\sigma^2]$, then there exists $\sigma^1\sigma \in \mathscr{L}$ with $[\sigma^1\sigma] = [\sigma^2]$.

Under (WE), we can still extend a feasible string by appending to it the extra events in another feasible string with a larger score. However, (WE) does not guarantee that the score space is max-closed.

The following consequence of (WE) is easily verified:

PROPOSITION 10.1. *If a scheme satisfies (WE), then it satisfies permutability and mutual interruption.*

To complement (WE), we need the following modification of an earlier boundary condition:

(B1′)(i) If α activates β, then $(\alpha, \beta) \in \mathscr{S}_-$;

(ii) If α and β interrupt each other, then $(\alpha, \beta) \in \mathscr{S}_+$.

We now have

LEMMA 10.2. *Lemma* 5.1 *holds for schemes that satisfy* (WE) *with condition* (B1) *replaced by* (B1′).

PROOF. Modification of the previous proof is required only to take account of the fact that the score space is no longer max-closed. This makes it possible for $\mathbf{d} + e_\alpha$ and $\mathbf{d} + e_\beta$ to be feasible even if $\mathbf{d} + e_\alpha + e_\beta$ is not. In this case, α and β interrupt each other. Hence $\bar{h}$ is (α, β)-supermodular, according to (B1′)(ii). The infinite cost penalty at $\mathbf{d} + e_\alpha + e_\beta$ ensures that $\bar{h}$ is also (α, β)-supermodular at $\mathbf{d}$.

Case (a) of the proof of Lemma 5.1 remains valid. Specifically, since $\mathbf{d}$ and $\mathbf{d} + e_\alpha + e_\beta$ are feasible scores, (WE) guarantes that σ^1 (with score $\mathbf{d}$) can be extended to either $\sigma^1\alpha\beta$ or $\sigma^1\beta\alpha$. Since $\sigma^1\beta\alpha$ is infeasible ($\beta \notin \mathscr{E}(\mathbf{d})$), $\sigma^1\alpha\beta$ must be feasible, implying that α is in $\mathscr{E}(\mathbf{d})$ and hence that α activates β. Case (b) remains intact. □

In previous sections, monotonicity of optimal controls followed from two basic steps: moving from a state-space formulation to a score-space formulation and showing that $\mathscr{S}$-modularity is preserved when costs are extended beyond boundaries. The first step followed from strong permutability and the fact that, under (SE), if $\mathbf{d}$ and $\mathbf{d} + e_\alpha$ are feasible then $\alpha \in \mathscr{E}(\mathbf{d})$. This property continues to hold if strong permutability is combined with (WE), so the transformation to scores remains valid. Lemma 10.2 verifies preservation of $\mathscr{S}$-modularity. Thus, the monotonicity results that follow from Lemma 5.1 generalize accordingly based on Lemma 10.2:

THEOREM 10.3. *Theorems* 5.2, 8.2(i–ii), *and* 8.5 *hold for schemes satisfying strong permutability and* (WE), *with condition* (B1) *replaced by* (B1′).

All of the earlier results referred to in Theorem 10.3 are based on extending costs beyond boundaries by using infinite penalties. Lemma 10.2 allows us to do the same under (WE) by strengthening (B1).

In principle, the projection method of §6 can also be used with (WE) and strengthened boundary conditions. A difficulty arises in selecting the appropriate projection. Under (WE), the score space need not be max-closed so an infeasible score may have several maximal feasible subscores; that is, it may have several bases rather than just one. Any rule for choosing among bases defines a basis extension. If it can be shown (as in Lemma 6.2) that the extended function is $\mathscr{S}$-modular, then the monotonicity established in Theorem 6.3 continues to hold. We have not, however, identified a general way of selecting a basis that ensures preservation of $\mathscr{S}$-modularity.

Hajek's (1984) model satisfies (WE) and strong permutability, even if switching between queues is incorporated. Hajek's state-space projection $(x_1, x_2) \mapsto (x_1^+, x_2^+)$ (also discussed in §6) implicitly defines a basis extension for functions on the score space that preserves $\mathscr{S}$-modularity. This model thus provides an instance in which the approach of §6 extends to a model satisfying weak exchange but not strong exchange.

10.2. *A loss system.* We conclude with a further example requiring the generalization from noninterruption to mutual interruption. We consider a *loss* system consisting of N identical serves but no room for waiting jobs. Jobs of K different classes arrive to use any available server. Such a system has been used to model a shared memory (Foschini and Gopinath 1983) and a link in a circuit-switched network (Kelly 1991), among other things. The required structural conditions extend to networks of such nodes as well.

Let $\alpha_1, \ldots, \alpha_K$ denote the arrival events for the K classes of jobs and let $\beta_1, \ldots, \beta_K$ denote the corresponding service-completion events. Since we consider only the Markov case, we may assume that the arrival events are inactive whenever the system is full. Similarly, we may assume that a total service effort μ_{β_i} is devoted to class i jobs without specifying how it is allocated among those jobs.

The state of the system is $x = (x_1, \ldots, x_K)$, where x_i is the number of type i jobs in service, $i = 1, \ldots, K$. Strong permutability is clearly satisfied: $x_i = \mathbf{d}_{\alpha_i} - \mathbf{d}_{\beta_i}$, for all i. A straightforward but somewhat tedious argument shows that (WE) is also satisfied. The score space,

$$\left\{\mathbf{d} \in Z_+^{2K} : \mathbf{d}_{\alpha_i} \geqslant \mathbf{d}_{\beta_i}, i = 1, \ldots, K;\ \sum_{i=1}^{K} \left(\mathbf{d}_{\alpha_i} - \mathbf{d}_{\beta_i}\right) \leqslant N\right\}$$

is neither max-closed nor min-closed.

The relations among the events are as follows:

(1) for $i = 1, \ldots, K$, α_i activates β_i when $x_i = 0$;

(2) for $i, j = 1, \ldots, K$, β_i activates α_j in any state x with $x_1 + \cdots + x_K = N$;

(3) for distinct $i, j = 1, \ldots, K$, α_i and α_j, $i \neq j$, are mutually interruptive in any state x with $x_1 + \cdots + x_K = N - 1$.

For simplicity, consider a separable cost function $h(x) = h_1(x_1) + \cdots + h_K(x_K)$. If each h_i is convex, then (B1′) holds. We thus have

PROPOSITION 10.4. *Suppose h_i, $i = 1, \ldots, K$, are convex and all rates are controllable to zero. Then there exists an optimal control μ such that each μ_{α_i} increases in every $\mathbf{d}_{\beta_j}$ and decreases in every $\mathbf{d}_{\alpha_k}$, $k \neq i$, and each μ_{β_i} increases in every $\mathbf{d}_{\alpha_j}$.*

Acknowledgement. Research of both authors was supported in part by NSF Grant MSS-92-16490. Research of David Yao was also supported in part by NSF Grant ECS-89-96230.

References

Bártroli, M. and Stidham, S., Jr. (1987). Towards a Unified Theory of Structure of Optimal Policies for Control of Networks of Queues. Working Paper, Department of Operations Research, University of North Carolina, Chapel Hill, NC.

Björner, A. (1985). On Matroids, Groups, and Exchange Languages. In *Matroid Theory and its Applications.* (L. Lovász and A. Recski Eds.), Coll. Math. Soc. J. Bolyai **40**, North-Holland, Amsterdam, pp. 25–60.

Chen, H., Yang, P., and Yao, D. D. (1991). Control and Schedling in a Two-Station Queueing Network: Optimal Policies and Heuristics. Working Paper, Faculty of Commerce, University of British Columbia.

Dietrich, B. L. (1989). Matroids and Antimatroids—A Survey. *Discrete Math.*, **78**, 233–237.

Foschini, G. J. and Gopinath, B. (1983). Sharing Memory Optimally, *IEEE Trans. Comm.* **31** 352–360.

Glasserman, P. (1991). Structural Conditions for Perturbation Analysis Derivative Estimation: Finite-Time Performance Indices. *Oper. Res.* **39** 724–738.

_____ and Yao, D. D. (1992a). Monotonicity in Generalized Semi-Markov Processes. *Math. Oper. Res.* **17** 1–21.

_____ and _____ (1992b). Generalized Semi-Markov Processes: Antimatroid Structure and Second-Order Properties. *Math. Oper. Res.* **17** 444–469.

Ghoneim, H. A. and Stidham, S., Jr., (1985). Control of Arrivals to Two Queues in Series. *European J. Oper. Res.* **21** 399–409.

Hajek, B. (1984). Optimal Control of Two Interacting Service Stations. *IEEE Trans. Automat. Control* **29** 491–499.

_____ (1985). Extremal Splittings of Point Processes. *Math. Oper. Res.* **10** 543–556.

Kelly, F. P. (1991). Loss Networks. *Ann. Appl. Probab.* **1** 319–378.

Korte, B. and Lovász, L. (1983). Structural Properties of Greedoids. *Combinatorica* **3** 359–374.

Kumar, P. R. (1993). Re-Entrant Lines. *QUESTA* **13** 87–110.

Melamed, B. (1982). On Markov Jump Processes Embedded at Jump Epochs and Their Queueing-Theoretic Applications. *Math. Oper. Res.* **7** 111–128.

Rosberg, Z., Varaiya, P. P., and Walrand, J. C. (1982). Optimal Control of Service in Tandem Queues, *IEEE Trans. Automat Control* **27** 600–609.

Schassberger, R. (1976). On the Equilibrium Distribution of a Class of Finite-State Generalized Semi-Markov Processes. *Math. Oper. Res.* **1** 395–406.

Serfozo, R. F. (1981). Optimal Control of Random Walks, Birth and Death Processes, and Queues. *Adv. in Appl. Probab.* **13** 61–83.

_____ (1989). Poisson Functionals of Markov Processes and Queueing Networks. *Adv. in Appl. Probab.* **21** 595–611.

Shor, P. W., Björner, A., and Lovász, L. (1991). Chip-Firing Games on Graphs. *European J. Combinatorics* **12** 283–291.

Topkis, D. M. (1978). Minimizing a Submodular Function on a Lattice. *Oper. Res.* **26** 305–321.

Veatch, M. H. and Wein, L. M. (1992). Monotone Control of Queueing and Production/Inventory Systems. *QUESTA* **12** 391–408.

Weber, R. R. and Stidham, S., Jr. (1987). Optimal Control of Service Rates in Networks of Queues. *Adv. in Appl. Probab.* **19** 202–218.

Whitt, W. (1980). Continuity of Generalized Semi-Markov Processes. *Math. Oper. Res.* **5** 494–501.

Whittle, P. (1983). *Optimization Over Time*, *Vol.* II. John Wiley & Sons, NY.

P. Glasserman: 403 Uris Hall, Columbia Business School, New York, New York 10027; e-mail: fapglass@cugsbvm.gsb.columbia.edu

D. D. Yao: IE/OR Department, Columbia University, New York, New York 10027; e-mail: yao@cunixd.cc.columbia.edu