

Stochastic Depletion Problems: Effective Myopic Policies for a Class of Dynamic Optimization Problems

Carri W. Chan

Department of Electrical Engineering, Stanford University, Stanford, California 94305,
cwchan@stanford.edu

Vivek F. Farias

Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02142, vivekf@mit.edu

This paper presents a general class of dynamic stochastic optimization problems we refer to as *stochastic depletion problems*. A number of challenging dynamic optimization problems of practical interest are stochastic depletion problems. Optimal solutions for such problems are difficult to obtain, both from a pragmatic computational perspective as well as from a theoretical perspective. As such, simple heuristics are desirable. We isolate two simple properties that, if satisfied by a problem within this class, guarantee that a *myopic* policy incurs a performance loss of at most 50% relative to the optimal adaptive control policy for that problem. We are able to verify that these two properties are satisfied for several interesting families of stochastic depletion problems and, as a consequence, we identify computationally efficient approximations to optimal control policies for a number of interesting dynamic stochastic optimization problems.

Key words: stochastic optimization; scheduling; approximations

MSC2000 subject classification: Primary: 93E20, 90C39; secondary: 90C59, 90B36, 90B22

OR/MS subject classification: Primary: optimal control; approximations; secondary: stochastic model applications

History: Received January 17, 2008; revised June 3, 2008 and August 15, 2008. Published online in *Articles in Advance* April 17, 2009.

1. Introduction. This paper presents a general class of dynamic stochastic optimization problems we refer to as *stochastic depletion problems*. Our study of this general class of problems is motivated by the fact that a number of dynamic optimization problems of practical interest are stochastic depletion problems. In fact, certain deterministic optimization problems that have been the focus of a good amount of recent research are also special cases of stochastic depletion problems.

Informally, a stochastic depletion problem is specified by *item types* and *activity sets*. The use of an activity results in the depletion of items of various types. In particular, the number of items of a particular type depleted at any point in time are randomly distributed according to a distribution specified by a set of parameters specific to the activity employed at that time and the number of items of that type available. The parameters specifying these distributions are themselves given by exogenous stochastic processes. New items of a given type may appear and existing items depart according to exogenous stochastic processes. Item depletion generates rewards, and activities must be selected adaptively over time to accomplish such depletion. An adaptive activity selection policy in this framework has knowledge of system dynamics and at any given time must select an activity. An optimal such policy generates maximum total expected reward. We will subsequently see that a number of difficult dynamic optimization problems ranging from optimal control of parallel server queueing models to optimal ad-words allocation can be cast as stochastic depletion problems. It is not surprising then, that finding an optimal control policy for a stochastic depletion problem is, in general, difficult.

We identify two simple properties we refer to as value function monotonicity (VFM) and the immediate rewards (IR) property that, if satisfied by a stochastic depletion problem, guarantee that a *myopic* policy generates expected reward within a factor of $1/2$ of the optimal adaptive policy for that problem; that is, the myopic policy is a *2-approximation* algorithm for that problem. This policy is allowed to use all the information available up to the current point in time and maximizes expected reward earned over the following time step. Both properties are intuitive: the VFM property states that the optimal total expected reward (or value) accrued in the future starting from a particular state of the system is nondecreasing in the vector of available items at that state. In the other direction, the IR property states that the additional value gained by making available additional items at a particular state of the system is at most the reward earned for the depletion of those items.

We are able to verify both the VFM and IR properties for large families of stochastic depletion problems. These include stochastic depletion problems for which the total reward for items depleted over time is given by a nondecreasing submodular function of the vector of items depleted. An available item of a given type is depleted independent of all other items in the system with a probability that depends on time, the item type, and the activity employed. We refer to these as submodular stochastic depletion problems. We are also able to address families of problems where the reward earned for depleting an item is a nonincreasing function of the time of

depletion, specific to that item's type. We refer to these as linear decaying stochastic depletion problems. Our performance analysis of the myopic heuristic is sharp for both these families of stochastic depletion problems.

Our systematic study of the general class of problems presented in this paper results in several contributions which we now outline.

Stochastic Control Problems: Our framework lets us easily recognize and analyze simple effective heuristic policies for several high-dimensional stochastic control problems for which finding optimal control policies is difficult.

For instance, control problems pertaining to several interesting discrete-time queueing models with general arrival processes and geometric service times may be reduced to stochastic depletion problems provided one allows for service disciplines with pre-emption. As an example, we consider a discrete-time equivalent of a well-studied parallel server queueing model that has been an important subject of recent research (see, for instance Harrison and Zeevi [16], Bassamboo et al. [3, 2]). Finding optimal service policies for such problems is typically challenging. We identify a simple myopic policy that is a constant factor approximation algorithm for a broad class of performance metrics. Our policy bears similarities to the so-called " $c - \mu$ " scheduling rules (see, for instance, van Mieghem [29]).

As another example we consider problems of dynamic "product line design." These are problems where a firm must dynamically adjust the assortment of products it offers for sale so as to maximize expected revenues. Sales of a given product are influenced by the entire assortment of products offered as well as the prevailing market sizes for various customer segments. In addition to generating revenues, sales in a given period impact market sizes in subsequent periods. This represents an important generalization of static product line design problems of the type considered by van Ryzin and Mahajan [30]. Our analysis yields an effective heuristic that simply calls for the solution of a sequence of *static* product line design problems that in effect account for system dynamics only to a limited extent.

Online algorithms for stochastic variants of well-studied deterministic problems: We show that a myopic policy earns expected rewards that are within a constant factor of the optimal adaptive policy for what we believe to be important online stochastic generalizations of a number of problems studied typically in deterministic settings. The stochastic generalizations we present go beyond what may be modeled in the traditional online versions of these problems and incorporate features we view as desirable from a modeling perspective. We present approximation guarantees for these generalizations that are typically no worse (and sometimes, better) than the best-known guarantees for their deterministic counterparts.

For instance, we are able to provide an efficient myopic policy that is a 2-approximation algorithm for a stochastic broadcast scheduling problem. Successful data transmission in our broadcast scheduling model is *stochastic* which makes it naturally applicable in several communications engineering contexts. With a bound on the maximal number of simultaneous broadcasts the best-known scheme for the *deterministic* broadcast scheduling problem (where page transmission is successful with probability 1) is an offline 4-approximation (Bar-Noy et al. [1]). With no constraints on the maximal number of simultaneous broadcasts, the best-known offline algorithm for the deterministic broadcast scheduling problem is a 4/3-approximation due to Gandhi et al. [14] and the best-known online algorithm for the same scenario is a 2-approximation due to Kim and Chwa [20]. Our contribution to this body of work is to demonstrate that even with stochasticity in data transmission one can still achieve good constant factor performance guarantees via an adaptive myopic algorithm.

As another example, we consider a stochastic generalization to the AdWords Assignment problem (Fleischer et al. [12]) where revenues are generated via stochastic clicks on placed ads as opposed to simply via the placement of an ad as in the deterministic version of the problem. Using an appropriate approximation algorithm as a subroutine for the myopic problem we demonstrate an efficient $(3 + \epsilon)$ myopic scheme which matches the best-known $(3 + \epsilon)$ -approximation algorithm for the deterministic version of the AdWords Assignment problem due to Goundan and Schulz [15].

A unified framework: We provide an elegant, unified framework for the design and analysis of stochastic optimization problems that is analogous to that for submodular maximization over simple matroids. In particular, submodular maximization problems over several simple matroids (such as the cardinality and partition matroid) can be captured as submodular stochastic depletion problems. A number of interesting problems (such as the AdWords Assignment problem) are known to be examples of such submodular optimization problems (see Goundan and Schulz [15], for example). The stochastic depletion framework provides a natural vehicle for the analysis of stochastic variants of such problems wherein the notion of selecting a set element translates to an attempt at selection; the success of such an attempt is specified by an exogenous stochastic process.

We anticipate that the characterization of dynamic stochastic optimization problems that admit simple control policies with constant factor guarantees provided in the present work, is likely to allow for the analysis of simple

heuristics for a number of problems beyond the handful of examples we have alluded to above. In particular, the two abstract properties that guarantee the effectiveness of a myopic policy are typically not hard to recognize and could potentially be established for families of problems outside those discussed here. The remainder of this paper is organized as follows: In §2 we formally specify the class of stochastic depletion problems. Section 3 presents a myopic heuristic for stochastic depletion problems and identifies two simple properties—the VFM and IR properties—that, if satisfied by a stochastic depletion problem, guarantee that a myopic policy is a 2-approximation algorithm for that problem. Section 4 verifies the VFM and IR properties for two general families of stochastic depletion problems—submodular stochastic depletion problems and linear decaying stochastic depletion problems while the following two sections discuss a number of applications that lie within these families. Section 7 concludes with a perspective on interesting directions for further work.

2. Model. We are given a collection of *items*, each of which belongs to one of M types indexed by m . There can be at most \bar{x}_m items of type m available at any time. Items are depleted via the execution of a suitable activity from a set of feasible activities \mathcal{A} , and depletion of a set of items garners a nonnegative reward we will formalize shortly. Time is discrete (indexed by $t \in [0, T]$) and in each time-step one must choose to employ some activity from \mathcal{A} ; we let T denote the length of the time horizon. We let i index the elements of \mathcal{A} and denote a general element of \mathcal{A} by A . Let $x_{t,m}$ denote the number of items of type m that remain at the start of the t th time-step. Assuming one chooses action $A \in \mathcal{A}$ in the t th time-step, the number of items of each type m depleted within that time-step is given by a $\times_{m=1}^M \{0, \dots, x_{t,m}\}$ -valued random vector, X_t^A . The distribution of the random vector X_t^A is specified by a parameter $P_t(A)$ which itself is given by an exogenous stochastic process $\{P_t(A)\}$, specified for all A and taking values in some compact set \mathcal{R} . X_t^A is assumed independent of the past and $P_{t'}(A')$ for all $A', t' > t$, given $P_t(A)$ and the vector of nondepleted items x_t . For example, we may have that $\{P_t(A)\}$ is a $[0, 1]^M$ -valued stochastic process and assuming one chooses action A at time t , $X_{t,m}^A$ is an independent Binomial $(x_{t,m}, P_{t,m}(A))$ random variable. In what follows we drop the superscript A from X_t^A for economy, as the dependence on A will be clear from context. We have $x_{t+1,m} = x_{t,m} - X_{t,m}$ for all m , and receive a total reward of $g(x_t, x_{t+1}, t)$ where $g: \mathbb{Z}_+^M \times \mathbb{Z}_+^M \times [0, T] \rightarrow \mathbb{R}_+$ satisfies:

ASSUMPTION 2.1. *For all $x, x' \in \mathbb{Z}_+^M$, $g(x, x', \cdot)$ is a nonincreasing, nonnegative function. In addition, we assume $g(x, x', T) = 0$ for all $x, x' \in \mathbb{Z}_+^M$.*

Our objective is to design an adaptive scheduling policy that maximizes total expected reward earned within the first T time-steps. We define as our *state-space* the set

$$\mathcal{S} = \left\{ (x, t, p_1, p_2, \dots, p_{|\mathcal{A}|}): x \in \times_m \{0, 1, \dots, \bar{x}_m\}, 0 \leq t \leq T, p_i \in \mathcal{R}^{t+1} \forall i \right\}.$$

In particular, a state is associated with a vector of items remaining to be depleted, time, and a history of the P_t processes. We denote by $x(s)$ the projection of s onto its first co-ordinate and similarly employ the notation $t(s)$ and $p_i(s)$ for $i = 1, 2, \dots, |\mathcal{A}|$. We let the random variable $S_t \in \mathcal{S}$ denote the state in the t th epoch. We digress briefly to note that if the P_t processes were assumed to be given a priori—a scenario that will become relevant to our theoretical analysis in what follows—then the description of the state space may be significantly simplified. In particular, we may in this case consider as our state space the set

$$\mathcal{S} = \left\{ (x, t): x \in \times_m \{0, 1, \dots, \bar{x}_m\}, 0 \leq t \leq T \right\},$$

where a state is associated with a vector of items remaining to be depleted and time.

Finally, we define the random reward function $R: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}_+$ according to $R(s, A) = g(x(s), x(s) - X_{t(s)}^A, t(s))$, where $X_{t(s)} = X_{t(s)}^A$. We note that since at time t the realization of X_t from taking a particular action is unknown, *any* control policy is a priori unaware of the exact reward accrued from a particular action; only the distribution of this reward is known. We define a control policy π as a mapping from \mathcal{S} to the set of feasible activities \mathcal{A} , and we denote by Π the set of all such policies. Define the expected total reward-to-go under a policy π starting at state s according to:

$$J^\pi(s) = E \left[\sum_{t'=t(s)}^{T-1} R(S_{t'}, \pi(S_{t'})) \mid S_{t(s)} = s \right].$$

We let $J^*(s) = \sup_{\pi \in \Pi} J^\pi(s)$ denote the maximum expected total reward-to-go under any policy. The supremum in the definition of J^* is achieved and we denote by π^* the corresponding optimal policy; i.e., $\pi^*(s) \in \arg \sup_{\pi \in \Pi} J^\pi(s)$. We will refer to the problem of finding such an optimal policy π^* as a *stochastic depletion problem*.

We remark that our formulation permits modeling exogenous item arrivals and deadlines on the latest permissible time of depletion for a given item. In particular, assuming without loss that $\bar{x}_m = 1$, that is, a given item type can have at most a single item (otherwise, we could simply refine the definition of a type), we associate with each type an arrival time τ_m and a deadline $d_m \geq \tau_m$. One may then assume that the P_t processes are such that $X_{t,m} = 0$ a.s. for all $A \in \mathcal{A}$ if $t \notin [\tau_m, d_m]$ to model the fact that item type m arrives at time τ_m and may not be depleted beyond time d_m ; see §5.2 for a concrete illustrative example. Such a formulation succinctly assumes a (known) bound on the total number of arrivals in any given period.

The optimal reward-to-go function (or *value* function) J^* and the optimal scheduling policy π^* can in principle be computed via dynamic programming: in particular, letting $S(s, A)$ denote the random next state encountered upon employing activity A in state s , define the dynamic programming operator \mathcal{H} according to:

$$(\mathcal{H}J)(s) = \max_{A \in \mathcal{A}} E[R(s, A) + J(S(s, A))], \quad (1)$$

for all $s \in \mathcal{S}$ with $t(s) \leq T - 1$. J^* may then be found as the solution to the Bellman equation $\mathcal{H}J = J$, with the boundary condition $J(s') = 0$ for all s' with $t(s') = T$. The optimal policy π^* may be found as the greedy maximizer with respect to J^* in (1). Of course, this approach is computationally intractable: even in the event that the P_t processes are known a priori, the state space (the set of all (x, t)) is exponentially large. As such, this makes solution of a general stochastic depletion problem pragmatically difficult.

In addition to the above informal description for why we might expect finding an optimal solution to be a difficult task, one may easily see that special classes of stochastic depletion problems are NP-hard. We consider one such class here for completeness: Consider “rational clairvoyant” stochastic depletion problems where the P_t sequences are rational-valued deterministic sequences in $[0, 1]^M$. $X_{t,m}^A$ is a Binomial $(x_{t,m}, P_{t,m}(A))$ random variable for all t, m, A , independent of the past and $X_{t,m'}^{A'}$ for all $A', m' \neq m$, given $x_{t,m}$. We assume $g(x_t, x_{t+1}, t) = \sum_m w_m (x_{t,m} - x_{t+1,m})$ where $w_m \geq 0$ for all m . The input to such a problem consists of $|\mathcal{A}|$ rational-valued sequences of length T , and $2M$ rational numbers representing the initial number of jobs of each type and the reward constants, w_m , for each type. One may then construct a polynomial time reduction from the set-cover decision problem (which is NP-complete) to the rational clairvoyant stochastic depletion problem. This is formalized in the following result:

THEOREM 2.1. *The rational clairvoyant stochastic depletion problem is NP-hard.*

PROOF. We reduce the set-cover decision problem to the rational clairvoyant stochastic depletion problem. An instance of the set-cover decision problem is specified by a ground set \mathcal{U} , a cover set $\mathcal{B} \subseteq 2^{\mathcal{U}}$, and an integer k (where $k \leq |\mathcal{U}|$ without loss) and we must decide whether a cover (that is a subset of \mathcal{B} whose union is a superset of \mathcal{U}) of size $\leq k$ exists. We reduce this question to the optimal solution of the following rational clairvoyant stochastic depletion instance: We consider a problem with $|\mathcal{U}|$ item-types, and assume we have a single item of each type. We let each set in \mathcal{B} correspond to a feasible activity in the sense that the use of that activity results in the depletion of all items in that set with probability 1 and the depletion of items outside that set with probability 0 in any time slot. We let the depletion of a single item result in unit reward, $w_m = 1$, and assume that the time horizon for scheduling is k . Assuming a polynomial time algorithm for rational clairvoyant stochastic depletion, the reduced problem would require time that is $\text{poly}(|\mathcal{U}|, |\mathcal{B}|, k) = O(\text{poly}(|\mathcal{U}|, |\mathcal{B}|))$. If the optimal solution to this instance of the stochastic depletion problem has total reward $|\mathcal{U}|$, we know that there exists a set-cover of size $\leq k$. Conversely, if there exists a set-cover of size $\leq k$, then there exists a depletion policy with total reward $|\mathcal{U}|$. Our reduction is thus many-one and polynomial in the size of the input. This completes the proof. \square

We remark that the above reduction can also be used to reduce an interesting *optimization* problem related to set-covering—namely that of maximum set-coverage where one may pick at most k elements of \mathcal{B} so as to cover as many elements of \mathcal{U} as possible—to a stochastic depletion instance. In fact, this reduction is a special class of a useful set of reductions we will explore in §6.

We have, in this section, introduced a general class of dynamic stochastic optimization problems that as we shall see in later sections admit a number of interesting applications. Computing optimal solutions for such problems is evidently hard; the next section will present and analyze a natural, simple-to-implement heuristic for such problems.

3. A myopic heuristic for stochastic depletion. A natural heuristic policy one may consider for a stochastic depletion problem is given by the myopic policy which in state s chooses an activity set A that maximizes expected reward earned over the following time-step. That is,

$$\pi^g(s) \in \arg \max_{A \in \mathcal{A}} E[R(s, A)].$$

Such a policy is adaptive but ignores the evolution of the system and the impact of the present choice of activity on rewards in future states. The set \mathcal{A} in the myopic problem above is potentially exponentially large. In many cases however, this set has an implicit polynomial-sized representation (for instance, \mathcal{A} may correspond to a set of matchings) and the myopic maximization problem is efficiently solved. We will later also address the case where the myopic maximization problem is difficult but one has access to an appropriate near-optimal oracle.

3.1. The myopic heuristic is an online 2-approximation algorithm: Sufficient conditions. Our objective in this section will be to identify stochastic depletion problems for which the myopic heuristic is guaranteed to be within a constant factor to optimal. In particular, we will identify stochastic depletion problems for which we will have, for any state $s \in \mathcal{S}$, $J^*(s) \leq 2J^{\pi^g}(s)$.

Note that the myopic heuristic does not utilize any information about the evolution of the P_t processes. In particular, the immediate reward incurred in choosing activity A at some state s is independent of the past and $P_{t'}(A')$ for all $A', t' > t(s)$, given $x(s)$ and $P_t(\cdot)$. We will thus assume that the P_t processes are a priori given sequences; the expected total reward earned under the myopic heuristic in this clairvoyant scenario is equal to the expected total reward earned under the myopic heuristic for the corresponding sample path of the P_t processes. We will compare the performance of the myopic heuristic to that of an optimal *clairvoyant* algorithm that knows the realizations of the P_t processes a priori. Since an optimal clairvoyant policy must dominate the optimal policy, it will suffice to demonstrate performance guarantees relative to the optimal clairvoyant policy. Recall that such an optimal clairvoyant policy may be computed over a reduced state-space:

$$\mathcal{P} = \left\{ (x, t): x \in \prod_m \{0, 1, \dots, \bar{x}_m\}, 0 \leq t \leq T \right\}.$$

In the sequel, we will only consider such clairvoyant optimal policies; any reference to an optimal policy or value function in the sequel will pertain to an optimal policy or value function for the clairvoyant problem. Comparing performance to a clairvoyant policy yields performance guarantees that are valid over individual sample paths of the P_t processes. In particular, our guarantees will imply that the myopic heuristic is a *2-competitive online algorithm* where the optimal scheme is allowed knowledge of entire sample-paths of the P_t processes but *does not know* the realization of X_t until time $t+1$ (which is somewhat different from the typical competitive analysis setting).

We now identify two properties that if satisfied by the optimal clairvoyant value function J^* , will imply our desired approximation guarantee.

PROPERTY 1 (VALUE FUNCTION MONOTONICITY). *Consider states s, s' satisfying $x(s) \geq x(s')$, $t(s) = t(s')$. The VFM property requires that $J^*(s) \geq J^*(s')$. In words, all else being equal, it is advantageous to start at a state with a greater number of items available.*

Before we describe the second property we find it convenient to introduce some notation. For $\alpha \in \mathbb{Z}_+^M$, define a mapping $\tilde{S}_\alpha: \mathcal{S} \rightarrow \mathcal{S}$ according to $\tilde{S}_\alpha(s) = s'$ with $t(s') = t(s)$ and $x(s')_m = (x(s)_m - \alpha_m)^+$ for all m . $\tilde{S}_\alpha(s)$ is thus the state obtained if one were permitted to employ some set of activities (which presumably resulted in $X_t = \alpha$) but without incurring the use of a time-step.

PROPERTY 2 (IMMEDIATE REWARDS). *For all $\bar{s} \in \mathcal{S}$ and $\alpha \in \times_m \{0, 1, \dots, x(\bar{s})_m\}$,*

$$J^*(\bar{s}) \leq g(x(\bar{s}), x(\bar{s}) - \alpha, t(\bar{s})) + J^*(\tilde{S}_\alpha(\bar{s})).$$

This property states that it is advantageous if one were able to deplete items *without incurring the use of a time-step*. In particular, if instead of starting at some state $\bar{s} \in \mathcal{S}$, one started at state $\tilde{S}_\alpha(\bar{s})$ and was, in addition, given reward for the depletion of α items, this property requires that the value of the second scenario be at least as large as the first.

The two properties we have developed thus far for the optimal value function J^* are essentially all we need to prove an approximation guarantee for the myopic heuristic. We now present the proof of our main approximation guarantee which assumes the VFM and IR properties.

We first introduce relevant notation: For $\alpha \in \mathbb{Z}_+^M$, define a mapping $\hat{S}_\alpha: \mathcal{S} \rightarrow \mathcal{S}$ that specifies the next state obtained if one employed an activity set in state s which resulted in $X_{t,m} = \alpha_m \wedge x(s)_m$ for all m . In particular, we define \hat{S}_α according to $\hat{S}_\alpha(s) = s'$ with $t(s') = t(s) + 1$ and $x(s')_m = (x(s)_m - \alpha_m)^+$ for all m .

THEOREM 3.1. *Assuming J^* satisfies Properties 1 and 2, we have for all $s \in \mathcal{S}$, $J^*(s) \leq 2J^{\pi^g}(s)$.*

PROOF. The proof proceeds by induction on the number of time-steps that remain in the horizon, $T - t(s)$. The claim is trivially true if $t(s) = T - 1$ since both the myopic and optimal policies coincide in this case. Consider a state s with $t(s) < T - 1$ and assume the claim is true for all states s' with $t(s') > t(s)$.

Now if $\pi^*(s) = \pi^g(s)$, then the next states encountered in both systems are identically distributed so that the induction hypothesis immediately yields the result for state s . Consider the case where $\pi^*(s) \neq \pi^g(s)$. Denote by $X_{t(s)}^*$ and $X_{t(s)}^g$ random vectors of depleted items in period $t(s)$ under optimal and myopic policies, respectively, at state s . Let $\mathbf{0}$ be an M -dimensional 0 vector. We have:

$$\begin{aligned}
 J^*(s | X_{t(s)}^*, X_{t(s)}^g) &= E[R(s, \pi^*(s)) | X_{t(s)}^*] + J^*(\hat{S}_{X_{t(s)}^*}(s)) \\
 &\leq E[R(s, \pi^*(s)) | X_{t(s)}^*] + J^*(\hat{S}_0(s)) \\
 &\leq E[R(s, \pi^*(s)) | X_{t(s)}^*] + g(x(s), x(s) - X_{t(s)}^g, t(s) + 1) + J^*(\tilde{S}_{X_{t(s)}^g}(\hat{S}_0(s))) \\
 &\leq E[R(s, \pi^*(s)) | X_{t(s)}^*] + g(x(s), x(s) - X_{t(s)}^g, t(s)) + J^*(\tilde{S}_{X_{t(s)}^g}(\hat{S}_0(s))) \\
 &= E[R(s, \pi^*(s)) | X_{t(s)}^*] + E[R(s, \pi^g(s)) | X_{t(s)}^g] + J^*(\tilde{S}_{X_{t(s)}^g}(\hat{S}_0(s))) \\
 &= E[R(s, \pi^*(s)) | X_{t(s)}^*] + E[R(s, \pi^g(s)) | X_{t(s)}^g] + J^*(\hat{S}_{X_{t(s)}^g}(s)) \\
 &\leq E[R(s, \pi^*(s)) | X_{t(s)}^*] + E[R(s, \pi^g(s)) | X_{t(s)}^g] + 2J^{\pi^g}(\hat{S}_{X_{t(s)}^g}(s)), \tag{2}
 \end{aligned}$$

where the first inequality follows from the assumed VFM property for J^* upon noting that $x(\hat{S}_{X_{t(s)}^*}(s)) \leq x(\hat{S}_0(s))$. The second inequality follows from the IR property assumed of J^* upon taking $\bar{s} = \hat{S}_0(s)$ and $\alpha = X_{t(s)}^g$. The third inequality follows from Assumption 2.1 since g was assumed nonincreasing in time. The third equality follows from the identity $\tilde{S}_{X_{t(s)}^g}(\hat{S}_0(s)) = \hat{S}_{X_{t(s)}^g}(s)$ which, in turn, is simply a consequence of the definitions of \tilde{S}_α and \hat{S}_α . The final inequality follows from the induction hypothesis.

Now,

$$J^{\pi^g}(s) = E[R(s, \pi^g(s)) + J^{\pi^g}(\hat{S}_{X_{t(s)}^g}(s))]$$

and $E[R(s, \pi^g(s))] \geq E[R(s, \pi^*(s))]$ by the definition of the myopic policy π^g so that taking expectations in (2), we have:

$$\begin{aligned}
 J^*(s) &= E[J^*(s | X_{t(s)}^*, X_{t(s)}^g)] \\
 &\leq E[R(s, \pi^*(s))] + E[R(s, \pi^g(s))] + 2E[J^{\pi^g}(\hat{S}_{X_{t(s)}^g}(s))] \\
 &\leq 2J^{\pi^g}(s).
 \end{aligned}$$

This concludes the proof. \square

3.1.1. Performance with an approximate myopic oracle. We will subsequently encounter a number of examples for which the set \mathcal{A} is exponentially large, but admits some implicit polynomial representation allowing for efficient solutions to the myopic problem

$$\max_{A \in \mathcal{A}} E[R(s, A)].$$

Sometimes, however, this problem may itself be difficult to solve. In such scenarios the use of an oracle that is an α -approximation to this subproblem is, in fact, a $(1 + \alpha)$ -approximation to the original stochastic depletion problem. In particular, assume $\pi^{\text{approx}}: \mathcal{S} \rightarrow \mathcal{A}$ satisfies

$$E[R(s, \pi^{\text{approx}}(s))] \geq \frac{1}{\alpha} \max_{A \in \mathcal{A}} E[R(s, A)],$$

for all $s \in \mathcal{S}$. One may then establish the following result whose proof is omitted but entirely analogous to Theorem 3.1 above:

THEOREM 3.2. *Assuming J^* satisfies Properties 1 and 2, we have $\forall s \in \mathcal{S}$, $J^*(s) \leq (1 + \alpha)J^{\pi^{\text{approx}}}(s)$.*

4. Families satisfying the VFM and IR properties. The previous section identified two abstract properties—namely, the VFM and IR properties, that if satisfied yield uniform performance loss guarantees for the myopic heuristic, via Theorems 3.1 and 3.2. These properties are, in general, difficult to check. We establish in this section two simple yet fairly general families of stochastic depletion problems that satisfy Properties 1 and 2 thereby guaranteeing that the myopic heuristic is a 2-approximation algorithm for those families. Although there may certainly be other families of problems satisfying the VFM and IR properties, the families we identify in this section accommodate a number of interesting applications which will be the focus of §5 and §6.

4.1. Submodular stochastic depletion problems. We consider problems for which $\{P_t(A)\}$ is a $[0, 1]^M$ -valued stochastic process for all $A \in \mathcal{A}$. Assuming one chooses action A at time t , $X_{t,m}$ is a Binomial $(x_{t,m}, P_{t,m}(A))$ random variable that, given $x_{t,m}$ and $P_{t,m}(A)$, is independent of the past, $X_{t,m'}$ for $m' \neq m$ and $P_{t'}(A')$ for all $A', t' > t$. We assume submodular rewards. In particular, we assume $g(x_t, x_{t+1}, t) = w(\bar{x} - x_{t+1}) - w(\bar{x} - x_t)$, where $w: \mathbb{Z}_+^M \rightarrow \mathbb{R}$ satisfies:

ASSUMPTION 4.1. $w: \mathbb{Z}_+^M \rightarrow \mathbb{R}$ satisfies:

- (i) (Monotonicity) $w(y) \geq w(y')$ for $y \geq y'$.
- (ii) (Submodularity) For $e \in \mathbb{Z}_+^M$, $w(y + e) - w(y) \leq w(y' + e) - w(y')$ if $y \geq y'$.

Such a class of functions clearly satisfies Assumption 2.1. We need to demonstrate the VFM and IR properties. Recall that we will consider a clairvoyant optimal algorithm that knows a priori the realizations of the sample paths of the P_t processes. We first demonstrate the IR property. It turns out that doing so requires only the monotonicity of w ; the submodularity of w is not required for this property to hold.

LEMMA 4.1 (IMMEDIATE REWARDS). *We have for submodular stochastic depletion problems, for all $\bar{s} \in \mathcal{S}$ and $\alpha \in \times_m \{0, 1, \dots, x(\bar{s})_m\}$,*

$$J^*(\bar{s}) \leq w(\bar{x} - x(\bar{s}) + \alpha) - w(\bar{x} - x(\bar{s})) + J^*(\tilde{S}_\alpha(\bar{s})).$$

PROOF. Consider using the optimal policy starting at state \bar{s} , and let S_T^* be the random state under this policy at the end of the time horizon (that is, at time T), so that:

$$J^*(\bar{s}) = E[w(\bar{x} - x(S_T^*))] - w(\bar{x} - x(\bar{s})), \tag{3}$$

where the expectation is over the randomness in the system—namely, the random item depletion defined by the P_t sequences and chosen activities. Similarly, let \tilde{S}_T^* be the random state under the optimal policy at the end of the time horizon upon starting in state $\tilde{S}_\alpha(\bar{s})$ and, as above, we note that

$$J^*(\tilde{S}_\alpha(\bar{s})) = E[w(\bar{x} - x(\tilde{S}_T^*))] - w(\bar{x} - x(\bar{s}) + \alpha).$$

Let us reconsider the optimal policy starting at state \bar{s} and, in particular, let us partition the initial set of items into a set of *fictitious* and *real* items; we assume that we begin with α_m fictitious items of type m and $x(\bar{s})_m - \alpha_m$ real items of type m . This partitioning serves purely as a labeling of items and does not impact the evolution of the system in any fashion. In particular, if at some point in time t , we have $x_{t,m}^f$ and $x_{t,m}^r$ fictitious and real items of type m respectively, then using activity set A results in the depletion of $X_{t,m}^f$ and $X_{t,m}^r$ fictitious and real items respectively, where $X_{t,m}^f$ is a Binomial $(x_{t,m}^f, P_{t,m}(A))$ random variable and $X_{t,m}^r$ is a Binomial $(x_{t,m}^r, P_{t,m}(A))$ random variable (so that $X_{t,m}^f + X_{t,m}^r \stackrel{d}{=} X_{t,m}$). We are left with $x_{t,m}^f - X_{t,m}^f$ and $x_{t,m}^r - X_{t,m}^r$ fictitious and real items respectively. Let $\sum_t X_{t,m}^f$ ($\sum_t X_{t,m}^r$) denote the number of fictitious (real) items of type m depleted at the end of the time horizon by the optimal policy starting in state \bar{s} .

We now make two critical observations:

OBSERVATION 1. We observe that $\sum_t X_{t,m}^f \leq \alpha_m$ for all m by construction.

OBSERVATION 2. Due to the fact that given s and a choice of activity, the depletion of a given item of type m at time t is independent of the past and the depletion of any other item in the system at that time, $\sum_t X_{t,m}^r$ may be viewed as the number of items of type m depleted under some induced randomized sub-optimal policy, say π' , starting at state $\tilde{S}_\alpha(\bar{s})$. This induced policy π' assumes in state $\tilde{S}_\alpha(\bar{s})$ the existence of an additional α_m items of each type m and simulates depletion of those items without garnering any reward for them. It operates like the optimal policy would but on this modified state. More specifically, letting $s_0 = \tilde{S}_\alpha(\bar{s})$ and $\alpha_0 = \alpha$, we have:

$$\pi'(s_0) = \pi^*((x(s_0) + \alpha_0, t(s_0))).$$

Defining $\alpha_1 = x(S((\alpha_0, t(s_0)), \pi'(s_0)))$ and $s_1 = S(s_0, \pi'(s_0))$,

$$\pi'(s_1) = \pi^*((x(s_1) + \alpha_1, t(s_1))).$$

In general, defining $\alpha_t = x(S((\alpha_{t-1}, t(s_{t-1})), \pi'(s_{t-1})))$ and $s_t = S(s_{t-1}, \pi'(s_{t-1}))$,

$$\pi'(s_t) = \pi^*((x(s_t) + \alpha_t, t(s_t))).$$

It is worth noting that $\alpha_0 - \alpha_t = \sum_{t'=\bar{s}}^t X_{t'}^f$ while $x(s_0) - x(s_t) = \sum_{t'=\bar{s}}^t X_{t'}^r$.

We consequently have:

$$\begin{aligned}
J^*(\tilde{S}_\alpha(\bar{s})) &\geq J^{\pi'}(\tilde{S}_\alpha(\bar{s})) \\
&= E\left[w\left(\bar{x} - x(\bar{s}) + \alpha + \sum_t X_t^r\right)\right] - w(\bar{x} - x(\bar{s}) + \alpha) \\
&\geq E\left[w\left(\bar{x} - x(\bar{s}) + \sum_t X_t^f + \sum_t X_t^r\right)\right] - w(\bar{x} - x(\bar{s}) + \alpha) \\
&= E[w(\bar{x} - x(S_T^*))] - w(\bar{x} - x(\bar{s}) + \alpha) \\
&= J^*(\bar{s}) + w(\bar{x} - x(\bar{s})) - w(\bar{x} - x(\bar{s}) + \alpha),
\end{aligned}$$

where the first inequality follows from the optimality of π^* among all nonanticipatory policies. The first equality follows from our definition of the policy π' in Observation 2 and from the definition of $\tilde{S}_\alpha(\bar{s})$. The second inequality follows from the monotonicity of the function w and Observation 1: $\sum_t X_t^f \leq \alpha$. The second equality is again by our construction of the X_t^f and X_t^r processes. The final equality follows from (3). This completes the proof. \square

While the IR property required only the monotonicity of w , the VFM property requires both the monotonicity of w and its submodularity. This result is intuitive: a controller that starts at state s may simply assume that it starts at state s' and track state evolution accordingly. Assuming submodular rewards, applying the optimal policy to this (incorrectly tracked) state trajectory guarantees the policy a total expected reward of at least $J^*(s')$, so the optimal policy must certainly do at least as well. The submodularity required is somewhat subtle, but it is simple to construct counterexamples in the absence of submodularity. We have:

LEMMA 4.2 (VALUE FUNCTION MONOTONICITY). *We have for submodular stochastic depletion problems, for all $s, s' \in \mathcal{S}$ s.t. $x(s) \geq x(s')$, $t(s) = t(s')$, $J^*(s) \geq J^*(s')$.*

PROOF. Consider a coupling of the systems starting at state s and s' wherein both systems witness identical sample paths for the item depletion processes defined by the P_t sequences. More precisely, assuming that at time t , the systems are in states s_t and s'_t respectively, then given P_t , the number of items depleted in both systems are coupled so that if $x(s_t) \geq x(s'_t)$, and we employ activity set A in both systems, then, for all m , the number of successfully depleted items of type m in the s_t system, $X_{t,m}$ ($\sim \text{Binomial}(x_{t,m}, P_{t,m}(A))$) and the number of successfully depleted items of type m in the s'_t system, $X'_{t,m}$ ($\sim \text{Binomial}(x'_{t,m}, P_{t,m}(A))$) satisfy $X_{t,m} = X'_{t,m} + Y_{t,m}$ where $Y_{t,m}$ is an independent Binomial ($x(s_t)_m - x(s'_t)_m, P_{t,m}(A)$) random variable. A symmetric situation must hold if $x(s'_t) \geq x(s_t)$.

Now assume that the system starting at s' uses an optimal policy whereas the system starting at state s mimics the actions of the s' system (call this policy $\bar{\pi}$). It is simple to see that $\bar{\pi}$ is an admissible nonanticipatory policy.

Under our coupling, we have at $t = t(s)$ that the number of items of type m depleted in the system starting at state s is greater than the number of items depleted in the system starting at state s' . That is, $X_{t,m} \geq X'_{t,m}$. It then follows that

$$\begin{aligned}
R(s, \bar{\pi}(s)) &= w(\bar{x} - x(s) + X_t) - w(\bar{x} - x(s)) \\
&\geq w(\bar{x} - x(s) + X'_t) - w(\bar{x} - x(s)) \\
&\geq w(\bar{x} - x(s') + X'_t) - w(\bar{x} - x(s')) \\
&= R(s', \pi^*(s')).
\end{aligned}$$

That is, the reward earned in the system starting at state s is higher than that in the system starting at state s' ; the first inequality above uses the monotonicity of w while the second inequality employs the submodularity of w . Now, in addition, by our coupling, both systems transition to states $S_{t(s)+1}$ and $S'_{t(s)+1}$, respectively, satisfying $x(S_{t(s)+1}) = x(s) - X_t = (x(s') - X'_t) + (x(s) - x(s') - Y_t) \geq x(s') - X'_t = x(S'_{t(s)+1})$, so that we may repeat the above argument for time $t(s) + 1$. Continuing in this fashion we see that in every time step the $\bar{\pi}$ controlled system starting at state s earns at least as large a reward as the π^* controlled system starting in state s' . Taking expectations over the random item depletions (i.e., the X_t and X'_t processes), we have $J^{\bar{\pi}}(s) \geq J^*(s')$. Since $J^*(s) \geq J^{\bar{\pi}}(s)$, we are done. \square

In light of Lemmas 4.1 and 4.2, Theorem 3.1 lets us conclude that the myopic heuristic is a 2-approximation algorithm for submodular stochastic depletion problems.

4.2. Linear decaying stochastic depletion problems. We consider here a family of stochastic depletion problems closely related to the family just considered that also admit the VFM and IR properties and are particularly useful for many applications. As before, we consider problems for which $\{P_t(A)\}$ is a $[0, 1]^M$ -valued stochastic process for all $A \in \mathcal{A}$. Assuming one chooses action A at time t , $X_{t,m}$ is a Binomial $(x_{t,m}, P_{t,m}(A))$ random variable that, given $x_{t,m}$ and $P_{t,m}(A)$, is independent of the past, $X_{t,m'}$ for $m' \neq m$, and $P_{t',m}(A')$ for all $A', t' > t$. We assume linear rewards that are nonincreasing in time. In particular, we assume $g(x_t, x_{t+1}, t) = \sum_m w_{m,t}(x_{t,m} - x_{t+1,m})$, where $w_{m,t}$ is a nonnegative, nonincreasing function of t for all m (for the special case where $w_{t,m} = w_m \geq 0$ for all t this is merely a special case of a submodular stochastic depletion problem model we have considered). We can verify the immediate rewards property for such systems via a proof found in the appendix that closely follows Lemma 4.1. Again, recall that the following results apply to the clairvoyant optimal policy.

LEMMA 4.3 (IMMEDIATE REWARDS). *We have for linear decaying stochastic depletion problems, for all $\bar{s} \in \mathcal{S}$ and $\alpha \in \times_m \{0, 1, \dots, x(\bar{s})_m\}$,*

$$J^*(\bar{s}) \leq \sum_m \alpha_m w_{t(\bar{s}),m} + J^*(\tilde{S}_\alpha(\bar{s})),$$

In addition, we may verify the VFM property. The proof of the following lemma is essentially identical to that of Lemma 4.2 and is omitted.

LEMMA 4.4 (VALUE FUNCTION MONOTONICITY). *We have for linear decaying stochastic depletion problems, for all $s, s' \in \mathcal{S}$ s.t. $x(s) \geq x(s')$, $t(s) = t(s')$, $J^*(s) \geq J^*(s')$.*

In light of Lemmas 4.3 and 4.4, Theorem 3.1 lets us conclude that the myopic heuristic is a 2-approximation algorithm for linear decaying stochastic depletion problems.

4.3. A worst case example. Having established the VFM and IR properties for the two families of stochastic depletion problems just discussed, we immediately have that the myopic policy has expected value within a factor of 1/2 of the optimal policy for problems from either family. This analysis is sharp. In particular, we now present a problem instance that is, in fact, a member of both problem families and for which we have that the optimal policy has expected value that is a factor of $2 - \epsilon$ better than the myopic policy; $\epsilon > 0$ can be made arbitrarily small.

EXAMPLE 4.1 (MYOPIC SUB-OPTIMALITY). Consider the case where $M = 2$ and $T = 2$, $g(x, x', t) = (x_1 - x'_1) + (1 - \epsilon)(x_2 - x'_2)$. Assume that $\bar{x}_1 = \bar{x}_2 = 1$ and that $x_{0,1} = x_{0,2} = 1$. Let $\mathcal{A} = \{1, 2\}$. The (deterministic) P_t processes are defined as:

$$\text{For } A = 1: \quad P_{0,1}(1) = 1, P_{0,2}(1) = 0, P_{1,1}(1) = 1, P_{1,2}(1) = 0.$$

$$\text{For } A = 2: \quad P_{0,1}(2) = 0, P_{0,2}(2) = 1, P_{1,1}(2) = 0, P_{1,2}(2) = 0.$$

In words, the item of type $m = 1$ may be depleted in either time-step via the use of $A = 1$, whereas the item of type $m = 2$ may be depleted only in the first time-step via $A = 2$. Only one of activity 1 or 2 may be employed within a given time-step. The myopic heuristic will first choose activity set $\{1\}$ (which earns a reward of 1 via the depletion of the $m = 1$ type item) over activity set $\{2\}$ (which earns a reward of $1 - \epsilon$ via the depletion of the $m = 2$ type job). Consequently, under the myopic heuristic, $x_{1,1} = 0$, $x_{1,2} = 1$, and the heuristic is unable to complete the one remaining job in the second time-step, earning a total reward of 1. An optimal schedule would first choose activity set $\{2\}$ (which earns a reward of $1 - \epsilon$ via the completion of the $m = 2$ type job). Consequently, under the optimal schedule, $x_{1,1} = 1$, $x_{1,2} = 0$, and the schedule is able to complete the one remaining job in the second time-step via the use of activity set $\{1\}$ earning a total reward of $2 - \epsilon$. We thus see that $J^*(s_0) = (2 - \epsilon)J^{\text{myopic}}(s_0)$ here.

5. Applications: Stochastic control. In the previous section, we presented two families of stochastic depletion problems for which the myopic heuristic is a 2-approximation algorithm. We now consider several problems of stochastic control that are easily seen to be members of these families. We thereby establish uniform performance guarantees for myopic policies for these stochastic control problems.

5.1. Service policies for simple queueing models. The following is a discrete time version of a parallel-server queueing model that has received a good deal of attention: We have I buffers and J servers. Each buffer sees a general discrete time arrival process with the restriction that a given buffer can see at most a single arrival in a given time slot. For example, each buffer i may see an independent Bernoulli(λ_i) arrival process. A given

server j may be used to service any single job in the system in a given time slot. In particular, should server j be used to service a job arriving to buffer i , the service time is assumed to be an independent geometric random variable with mean $\mu_{i,j}$ (possibly ∞). We allow for pre-emption in our service discipline. While allowing for pre-emption is unrealistic in some scenarios—for instance, in a call-center—it is quite realistic for many other applications of the parallel server model such as parallel or distributed computing. Consider the following natural objective: completion of a job that arrives at buffer i earns a nonnegative reward $r_{i,d}$, where d is the time that job has remained in the system (that is, the delay experienced by that job). We assume $r_{i,d}$ is nonincreasing in d . At every point in time one must decide on a matching between servers and available jobs with a view to maximizing the expected reward earned over T periods.

It is not difficult to see that the above problem is an example of a linear decaying stochastic depletion problem. In particular, we define an item type for every tuple (i, τ) , where $i = 1, 2, \dots, I$ and $\tau = 0, 1, \dots, T - 1$. Thus an item type m is associated with an arrival buffer i_m and an arrival time τ_m . We can have at most a single item of a given type; i.e., $\bar{x}_m = 1$. The set of feasible activities \mathcal{A} is simply the set of all matchings of servers to item types. Given a particular matching, the probability of depletion for a given item type (or job), is determined by the server matched to that job or 0 if no server is matched to it. Of course, a job may not be depleted prior to arrival. In particular, we have for item type $m = (i_m, \tau_m)$,

$$P_{t,m}(A) = \mathbf{1}_{t \geq \tau_m} \left(\sum_j \mathbf{1}_{(i,j) \in A} 1 / \mu_{i,j} \right).$$

We define our reward function g according to $g(x_t, x_{t+1}, t) = \sum_m w_{m,t}(x_{t,m} - x_{t+1,m})$, where we assume $w_{m,t} = r_{i_m, (t-\tau_m)^+}$. In particular, the reward generated in the t th time step is given by $\sum_{i=1}^I \sum_{t'=0}^t r_{i,t-t'} X_{i,t-t'}$, where $X_{i,t-t'} = 1$ if a job arriving to buffer i at the start of time-step t' was completed at time t , and is 0 otherwise. Since both the VFM and IR properties hold for this family of stochastic depletion problems, we have via Lemmas 4.3 and 4.4, and Theorem 3.1 that the myopic policy generates total expected rewards that are within a factor of 1/2 of the optimal policy. In fact, we have shown that this performance guarantee holds relative to an optimal policy that has full knowledge of the entire job-arrival process!

Continuous time variants of the problem above have been the subject of much study and results on optimal control policies in various asymptotic parameter regimes are available (for example, see Bassamboo et al. [3]); the formulation we have discussed focuses on a different objective and a nonasymptotic parameter regime, thus complementing that body of work. Although distinct, it is interesting to note the similarity of our myopic policy with the so-called $c - \mu$ scheduling rules (see for instance, van Mieghem [29]) for scheduling jobs arriving to multiple buffers served by a single server with a view to minimizing total delay cost (every job incurs a buffer-dependent, typically linear, delay cost). In another example, still closer to our work, Buyukkoc et al. [4] consider scheduling jobs arriving to multiple buffers served by a single server in discrete time and show the $c - \mu$ rule to be optimal for arbitrary arrival processes and geometric service times if the service discipline permits pre-emption. Relative to Buyukkoc et al. [4], we consider in our model the situation with multiple servers, each of which can serve some subset of buffers. However, we maximize reward as opposed to minimizing cost. Rewards decrease with delay incurred in a piecewise linear fashion, eventually decreasing to zero; our model cannot accommodate negative rewards. In contrast, in the $c - \mu$ setting (necessarily nonnegative) costs increase linearly with delay.

5.2. Stochastic broadcast scheduling. We consider a broadcast communication system where a single data item may be simultaneously transmitted to multiple users. In particular, we consider the following problem: we have a set of U users (indexed by u) and a finite set of data items or “pages,” $P = \{1, \dots, n\}$, indexed by i . In every time slot $\tau \in \{1, \dots, T - 1\}$, any given user may generate a request for some page (or pages) he has not requested in the past. We assume that every request for a page is associated with a deadline $d \in \{1, \dots, T - 1\}$. Should a request for page i by user u be successfully satisfied prior to its deadline, the transmitter earns nonnegative reward r_i^u . We assume that the arrival process governing requests from users, as well as the deadlines associated with those requests, are exogenous stochastic processes and further assume a (known) bound on the number of requests that may arrive in any given time slot. In each time slot, a single page can be transmitted (although in what follows we could also consider allowing up to b pages). Due to the broadcast nature of the system, this page may be transmitted simultaneously to up to $k \geq 1$ users. The communication channel to users is stochastic, so that should a page be transmitted to a particular user u in time t , that user receives the page with some channel-dependent probability P_t^u which is itself an exogenous stochastic process. In each time slot, one must decide which page to transmit and to which k users to maximize the expected reward accrued over T time slots.

Approximation algorithms for *deterministic* broadcast scheduling (where transmissions are successful with probability 1 so that $P_t^u = 1$ for all t, u) have received quite a bit of attention. The best-known approximation algorithm is a 4-approximation due to Bar-Noy et al. [1]. Without any constraints on the number of requests that can be satisfied by a single broadcast (i.e., $k = \infty$), the best-known algorithm is a 4/3-approximation algorithm due to Gandhi et al. [14]. The best-known online algorithm for the same is a 2-approximation due to Kim and Chwa [20]. Specializing to this deterministic case, our myopic online algorithm improves upon the offline results for “finite batching” in Bar-Noy et al. [1] albeit for uniform item sizes. Modeling stochasticity in communication channels to users is important since in real-world systems, congestion and various physical phenomena cause significant uncertainty in the successful transmission of pages. Scheduling communications over stochastic channels is, of course, the focus of a substantial body of work in communications engineering. See, for instance, Eryilmaz et al. [10], Su and Tassiulas [27], Ren et al. [25] for models closely related to the broadcast scheduling model we have presented. Most of that body of work is either simulation driven or focuses on coarser performance metrics (such as throughput optimality). A recent push in packetized multimedia scheduling has examined the case of utility maximization in lossy networks (see Chou and Miao [6] and related work). In this stream of work, each packet is associated with a weight that represents the reduction in distortion achieved with the reception of that packet, much like our formulation. Many of these formulations call for the solution of intractable optimization problems and heuristics such as the greedy heuristic, or heuristics with limited look-ahead, become essential.

The stochastic broadcast scheduling problem we have presented may be cast as a linear decaying stochastic depletion problem. Every request is associated with four parameters (u, i, τ, d) representing the user, page, time of request, and request deadline, respectively. We associate an item type with each such request. Thus an item type m is identified by a request by user u_m for page i_m , with an arrival time, τ_m , and deadline, d_m . An activity $A \in \mathcal{A}$ is simply an assignment of a given page to k users. Given a particular choice of activity A , the probability of depletion of a particular item type is simply given by the quality of the channel to the user corresponding to that type, P_t^u , provided that user u is served under activity A ; else it is 0. Of course a request may not be satisfied prior to its arrival or following the expiration of its deadline, so that for item type $m = (u_m, i_m, \tau_m, d_m)$:

$$P_{t,m}(A) = \mathbf{1}_{d_m > t \geq \tau_m} \mathbf{1}_{(i_m, u_m) \in A} P_t^{u_m}.$$

We define our reward function g according to $g(x_t, x_{t+1}, t) = \sum_m w_m (x_{t,m} - x_{t+1,m})$, where we assume $w_m = r_{i_m}^{u_m}$. This is a linear decaying stochastic depletion problem and since both the VFM and IR properties hold for this family of stochastic depletion problems, we have via Theorem 3.1 that the myopic policy generates expected value that is within a factor of 1/2 of the optimal policy. Moreover, the complexity of the heuristic is essentially the lowest we may expect for this problem; i.e., computational complexity grows linearly with the number of requests to be considered and the horizon. This is also substantially lower than some of the tailored approximation methods (such as Gandhi et al. [14]) that require solving linear programs with a number of variables that grows linearly with the number of requests and the length of the horizon.

It is interesting to consider a special case of the stochastic broadcast scheduling problem we have presented: In particular, assuming that all requests are known at time $t = 0$ and further, that all these requests have deadline T , the myopic policy is, in fact, optimal if the channel to each user is “static,” i.e., $P_t^u = C_u$, for all u and if at most a single packet may be broadcast to all (or up to k) users in any time step. This is established via an interchange argument; the proof of the following lemma may be found in the appendix. The result for the case where $k < \infty$ may be established via an essentially identical (but notationally cumbersome) interchange argument.

LEMMA 5.1. *Suppose all requests are available at $t = 0$ and share deadline T . Then, if P_t^u is constant ($= C_u$) for all u , and at most a single packet may be broadcast to all users in a given time step (i.e., $b = 1$ and $k = \infty$), then $J^{\pi^s}(s_0) = J^*(s_0)$ for all $s_0 \in \mathcal{L}$.*

Lemma 5.1 allows one to interpret the myopic heuristic for the general stochastic broadcast scheduling problem as one that at every time-step t makes the simplifying assumption that all channels are *static* with success probabilities given by P_t^u and that no further arrivals will be observed. This is, in fact, a common engineering design principle for scheduling over dynamic channels. For instance, Tsibonis and Geogiadis [28], Dua and Bambos [7, 9], Huang et al. [18], Dua and Bambos [8], and Chou and Miao [6] all derive optimal scheduling policies for problems similar to the broadcast scheduling problem here under the assumption of a static channel and other simplifying assumptions. The hope is that in conjunction with frequent channel state re-estimation (that is, frequent re-estimation of channel success probabilities), the use of scheduling schemes so derived may prove to be a very effective heuristic. In addition to being simple to implement and typically fast in practice,

such an approach is robust to errors in specifying channel dynamics. Lemma 5.1 and Theorem 3.1 thus lend theoretical support to this popular design principle. In particular, one may simply design a scheduling scheme assuming a static channel; one then employs this scheme in tandem with repeated channel re-estimation. Put another way, simply accounting for channel *state* suffices to obtain levels of performance within 50% to optimal.

5.3. Dynamic product-line design. Consider a firm that is capable of producing an array of related products that may potentially be sold to one or more customer segments, each distinguished by its willingness to pay for various product features. For a variety of reasons (manufacturing capacity and cost, marketing capabilities, etc.), the firm may be constrained in the number of different products it is capable of simultaneously offering for sale. Further, external competition may impose limitations on the prices the firm can post for a given product. Faced with these restrictions, the firm must decide on a product line to offer with a view to maximizing revenues. This is the essence of product-line “design” problems that have been extensively considered in the operations research and marketing literature. For instance, the classic third-degree price-discrimination model of Pigou [24] forms the basis of design principles that center on explicit market segmentation (see Frank et al. [13]). Alternatively, assuming a model of customer preference for various product attributes, one may consider optimizing the attributes of products offered for sale to maximize revenues; customers “self-select” product types that are of greatest appeal in this case. A number of product-line design problems of this type have been considered in past literature; Moorthy [22], Kohli and Sukumar [21], van Ryzin and Mahajan [30], and Hopp and Xu [17] are a few examples. A common thread to this work, however, is their consideration of static models. In reality, demand shocks and demand seasonality make the optimal product-line design problem an inherently dynamic one. For instance, consider the following example that illustrates the potential importance of accounting for seasonality in demand.

EXAMPLE 5.1. A firm may offer at most one of two products (“outdated” or “new”) for sale at any epoch (over two successive sales epochs) to two distinct consumer segments—“bargain hunters” and “early adopters.” Bargain hunters will purchase only the outdated product for one dollar in the first period with probability 1 and will make no purchase in period 2. “Early adopters” will purchase only the new product with probability 1 in either epoch for $1 + \epsilon$ dollars. Assume we begin with an equal number of consumers in both segments. It is clear that a product-line selection strategy that accounts for seasonality (by delaying the introduction of the new product to the second period) will earn about twice the revenues earned by a myopic strategy over two sales epochs.

How does one interpret the above example and to what extent must one account for seasonality and demand shocks in designing product lines? Motivated by this question, we consider the following dynamic product-line design problem: a firm is capable of offering products from some set \mathcal{P} and must at any point in time offer a subset of products $A \subset \mathcal{P}$ with $|A| \leq k$. The firm’s products are purchased by I consumer segments and we let $x_{t,i}$ denote the size of the i th segment in the t th sales epoch. Assuming that the product line offered at time t is A , any segment i consumer present in the market at that time t will purchase a product in A with known probability $P_{t,i}(A)$. Such a sale garners the seller revenue p_i and the consumer is lost to the system so that the size of the i th segment after sales in that epoch become $x_{t,i} - X_{t,i}^A$, where $X_{t,i}^A$ is a Binomial $(x_{t,i}, P_{t,i}(A))$ random variable. We assume that the firm has modeled the dynamics of the P_t processes and wishes to maximize expected revenues over T sales epochs.

It is worthwhile discussing some of the salient features of our model.

(i) We model the fact that a customer from a given segment may satisfy his requirement by the purchase of one of several product types (a purchasing customer may substitute between products in A), and that the probability of a particular segment purchasing a product is influenced by the entire array of product types offered ($P_{t,i}$ is a function of A). Such a model permits customer self-selection but under the restriction that all substitutes (all products in \mathcal{P}) are offered at the same price to every segment; this is precisely the type of model considered by van Ryzin and Mahajan [30]. Alternatively, one may view the model as assuming that the seller has a means of directly segmenting customers (as in Frank et al. [13]) and allowing for segment-specific prices, i.e., prices p_i that depend on the segment i .

(ii) We allow for general models of demand seasonality. In particular, we make no assumptions on the dynamics of the P_t processes. Further, we explicitly model the impact of current sales on future demands (“market saturation”). In particular, the number of sales in a given epoch are influenced by the sizes of the various customer segments, $x_{t,i}$, in that epoch; sales in a given period cause a reduction in these sizes thereby lowering potential sales to those segments in subsequent periods.

(iii) We assume that the product-line designer has available an estimate of market size within each consumer segment; i.e., he can estimate $x_{t,i}$ over time. Such an assumption is potentially valid in several industries; see Yunes et al. [31] for a practical discussion of this issue.

(iv) Prices for each consumer segment are fixed. That is, a sale to a segment i customer must be at a price p_i that is independent of the product variant purchased. In reality, fixed prices may arise, for example, due to the need to align with prices offered by competitors.

(v) Finally, we assume no limitations to inventories of a given product so that all realized demand in a period can be met. This is a simplification relative to models such as those considered by van Ryzin and Mahajan [30]. Nonetheless, for high-margin products where the cost of a lost sale is far higher than that of an unsold unit of inventory, this is likely to be a relatively mild simplification.

It is simple to cast the above model as a linear decaying stochastic depletion problem. In particular we associate with each customer segment i an item type m . Our set of activities $\mathcal{A} = \{A: A \in \mathcal{P}, |A| \leq k\}$, and depletion probabilities for item type m are specified according to $P_{t,m}(A) = P_{t,i_m}(A)$. We define our reward function g according to $g(x_t, x_{t+1}, t) = \sum_m w_{m,t}(x_{t,m} - x_{t+1,m})$, where we assume $w_{m,t} = p_{i_m}$. This is a linear decaying stochastic depletion problem, so that Lemmas 4.3 and 4.4 with Theorem 3.1 immediately tell us that a myopic policy generates expected revenues within a factor of $1/2$ of the optimal policy.

From a managerial perspective, this suggests a robust recipe for dealing with demand shocks and seasonality: at every opportunity for product line update, one simply solves a static product-line design problem with suitably revised estimates of the relevant customer demand model and market sizes. That is, optimization that entirely ignores contingencies for future demand shocks or seasonality but accounts for the current demand environment is already likely to be quite good and, in particular, will provide expected revenues at least within 50% of an optimal scheme in the context of the dynamic product-line design model considered here.

6. Applications: Stochastic variants of submodular maximization problems over matroids. In this section, we turn our attention to the use of the stochastic depletion framework as a useful *stochastic* analogue to submodular maximization problems over simple matroids such as the cardinality matroid and the partition matroid. A number of hard deterministic optimization problems can frequently be reduced to problems of this nature and in doing so, finding good approximation algorithms for these problems is reduced to the task of finding a good oracle for the myopic subproblem (for a number of recently considered problems of this type, see, for instance, Goundan and Schulz [15]). Our hope is to produce good approximation algorithms for useful stochastic variants of such problems. As an illustration, we will later consider an important stochastic generalization of the AdWords Assignment problem considered by Fleischer et al. [12] and Goundan and Schulz [15].

Given a set E , let $\mathcal{U} = 2^E$. A *cardinality matroid* is a subset of \mathcal{U} of the type $\mathcal{M} = \{F \subset E: |F| \leq k\}$, where k is an integer. A *partition matroid* is a subset of \mathcal{U} of the type $\mathcal{M} = \{F \subset E: |F \cap E_i| \leq k_i \forall i\}$, where we assume $E = \bigcup_{i=0}^n E_i$ and that integers k_i for $i = 0, 1, \dots, n$ are given.

Consider optimization problems of the form

$$\max_{A \in \mathcal{M}} f(A), \tag{4}$$

where $f: 2^E \rightarrow \mathbb{R}^+$ is a nondecreasing, submodular function. A number of interesting combinatorial optimization problems are reduced to such maximization problems, where \mathcal{M} is a cardinality or partition matroid. We begin with establishing how such deterministic optimization problems are captured within the stochastic depletion framework.

\mathcal{M} is a cardinality matroid: We reduce (4) to a submodular stochastic depletion instance assuming \mathcal{M} is a cardinality matroid: We are given $M = |E|$ item types and assume that we begin with a single item of each type; i.e., $\bar{x}_m = x_{0,m} = 1$. Let $\mathcal{A} = \{1, 2, \dots, |E|\}$, where $P_{t,m}(j) = 1$ if $m = j$ and 0 otherwise. That is, we define an item type for each element of E and in every time-step we are allowed to deplete at most one item. We select as our reward function $g(x_t, x_{t+1}, t) = f(\bar{x} - x_{t+1}) - f(\bar{x} - x_t)$ and set $T = k$. Observe that the value of an optimal solution to this problem is precisely $J^*(\bar{x}, 0)$. With Lemmas 4.1 and 4.2, Theorem 3.1 then immediately yields:

COROLLARY 6.1. *The myopic heuristic is a 2-approximation algorithm for maximizing a nondecreasing submodular function f over a cardinality matroid.*

We remark that this is a weaker result than the well-known optimal approximation ratio of $e/(e-1)$ due to Nemhauser et al. [23]. The analysis of Theorem 3.1 applies to a far broader class of problems, and in light of Example 4.1, we cannot expect a tighter guarantee for the greedy heuristic via that general line of analysis.

\mathcal{M} is a partition matroid: We reduce (4) to a submodular stochastic depletion instance assuming \mathcal{M} is a partition matroid: We are given $M = |E|$ item types and set $\mathcal{A} = \{1, 2, \dots, |E|\}$. We index the elements of E by m which identifies a particular element of E with a particular item type and assume that the first $|E_0|$ elements correspond to the elements of E_0 , the next $|E_1|$ elements to the elements of E_1 , and so forth. We set the

time horizon $T = \sum_i k_i$ and define $n + 1$ partitions of this horizon according to $T_j = [\sum_{l=0}^{j-1} k_l, \sum_{l=0}^j k_l - 1]$. We assume $P_{t,m}(j) = 1$ iff $m = j$ and $t \in \{T_j: m \in T_j\}$. We select as our reward function $g(x_t, x_{t+1}, t) = f(\bar{x} - x_{t+1}) - f(\bar{x} - x_t)$. In words, we define an item type for each element of E and identify each subset E_i with a partition of time. At any point in time $t \in T_i$, we are allowed to deplete at most one available item from the partition E_i . Observe that the value of an optimal solution to this problem is precisely $J^*(\bar{x}, 0)$. The myopic heuristic for this stochastic depletion problem corresponds precisely to the “local greedy” heuristic introduced by Fisher et al. [11], and we re-capture their result, namely:

COROLLARY 6.2. *The local greedy heuristic is a 2-approximation for maximizing a nondecreasing submodular function f over a partition matroid.*

Both classes of problems alluded to above have natural stochastic generalizations. As a simple example, one may consider a stochastic generalization to the problem of submodular maximization over a cardinality matroid which we refer to as the “stochastic selection problem.” As opposed to selecting at most k elements from E , one is allowed k attempts at selecting elements of E . If at the t th selection attempt one attempts to select element $e \in E$, the attempt is successful with probability P_t^e , where $\{P_t^e\}$ is an arbitrary $[0, 1]$ -valued sequence specified for every $e \in E$. We would like to find an adaptive item selection policy that maximizes the expected value of the set of successfully selected items. It is easy to see that the stochastic selection problem includes as special cases appropriate stochastic generalizations of problems such as the maximum coverage problem. The problem of adaptively selecting items to maximize the expected value of the set of successfully selected items is seen to be a submodular stochastic depletion problem using precisely the reduction for the cardinality matroid above and one immediately has the following result.

LEMMA 6.1. *The myopic heuristic is a 2-approximation for the stochastic selection problem.*

As an aside, we note that if, in addition, one assumes that $P_t^e = C$, a constant for all t and e , it is simple to demonstrate that the myopic heuristic is, in fact, an $e/(e - 1)$ -approximation. This may be demonstrated as a corollary to the original result of Nemhauser et al. [23]: one simply considers coupling the optimal and myopic schemes so that on each sample path, both schemes have an identical number of successful placements.

We now consider in some detail a practically relevant stochastic generalization of the AdWords Assignment problem (Fleischer et al. [12]). The deterministic problem may be reduced to the maximization of a submodular function over a partition matroid (see Goundan and Schulz [15]). We reduce our stochastic generalization to a submodular stochastic depletion problem.

6.1. Cost-per-click AdWords Assignment. Consider the following optimization problem faced by firms that serve ads on the internet. We are given a set of N advertisers (indexed by i) and K keywords (indexed by k). The i th advertiser has a budget B_i (≥ 0) and submits to the firm a valuation $v_{i,k}$ for every keyword k . In every one of T periods, a keyword from the set of K keywords arrives according to some exogenous stochastic process. We assume that at most C advertisers’ ads can be assigned to the arriving keyword. We denote by k_t the index of the keyword arriving at time t . Should an advertiser i be assigned to an arriving keyword, k_t , at time t , and if, in addition, his ad is clicked on, he pays the firm the minimum of v_{i,k_t} and his remaining budget at time t ; this payment is subtracted from his available budget. If the ad is not clicked on, then no payments are made. We assume that should advertiser i be assigned to keyword k_t at time t , his ad is clicked on with probability P_t^{i,k_t} . Letting $V_{t,i}$ denote the random payment thus made by an advertiser in the t th period, we are interested in devising an adaptive ad-to-keyword assignment scheme that maximizes $E[\sum_{i=1}^N (B_i \wedge \sum_t V_{t,i})]$, that is, the expected revenues earned by the firm.

The above problem was considered in a deterministic offline setting by Fleischer et al. [12] where it was assumed that $P_t^{i,k} = 1$ for all t, i, k and, in addition, the sequence $\{k_t\}$ of arriving keywords was specified a priori; the variant we consider here is an important generalization of that model since in practice advertisers make payments only if their displayed ads are clicked on, which happens with some positive, but small, probability. In addition, our formulation also allows us to capture exogenous advertiser arrivals and departures from the system—in particular, we simply assume $P_t^{i,k} = 0$ for all times t prior to a customer’s arrival to the system and following his departure from the system.

This problem is easily cast as a submodular stochastic depletion problem: In particular, we define an item type m for every advertiser-keyword-time triple (i_m, k_m, t_m) , and assume a single item of each type, i.e., $x_{0,m} = \bar{x}_m = 1$, for all m . The set of feasible activities, \mathcal{A} , is the set of all subsets of item types such that each subset has cardinality at most C and contains at most one type specific to a given user i . The probability that an item

of type m is depleted at time t assuming one selects activity A is given by

$$P_{t,m}(A) = \mathbf{1}_{\{(i_m, k_m, t_m) \in A, t=t_m, k_t=k_m\}} P_t^{i_m, k_m}.$$

Finally, the reward function, $g(x_t, x_{t+1}, t) = w(\bar{x} - x_{t+1}) - w(\bar{x} - x_t)$, where $w: \mathbb{Z}^M \rightarrow \mathbb{R}^+$ is defined according to

$$w(x) = \sum_i \left(B_i \wedge \sum_{m: i_m=i} v_{i_m, k_m} x_m \right)$$

and thus satisfies Assumption 4.1. We finally note that the myopic subproblem is trivial; it corresponds to choosing the C highest expected revenue advertisers so that our heuristic is essentially the computationally simplest heuristic one may consider for this problem with complexity that grows linearly in the number of keywords and ads served over the horizon. With Lemmas 4.2 and 4.1, Theorem 3.1 yields:

COROLLARY 6.3. *The myopic heuristic is a 2-approximation to the cost-per-click AdWords Assignment.*

In our formulation, a feasible ads-to-keyword assignment was subject to a simple cardinality constraint: an arriving keyword could have at most C ads assigned to it. We could instead consider using other, more complex constraints: in particular, in the formulation of Fleischer et al. [12], every ad is associated with a rectangle of a specific height and width, and every arriving keyword is associated with an available rectangular display area; a feasible assignment of ads-to-keywords is determined by a feasible packing of ad rectangles within the display rectangle. Using the max-weight rectangle packing $(2 + \epsilon)$ -approximation algorithm of Jansen and Zhang [19] for the myopic subproblem yields via Theorem 3.2, a $(3 + \epsilon)$ -approximation guarantee for the myopic heuristic which matches the best-known approximation guarantee available for the original deterministic problem (see Goundan and Schulz [15]). In the deterministic case, our heuristic coincides precisely with that of Goundan and Schulz [15]; the stochastic generalization entails essentially no increase in computational complexity.

7. Concluding remarks. We have in the present work introduced a general class of dynamic stochastic optimization problems—stochastic depletion problems. We believe this to be an interesting class of problems: in spite of being fairly general, stochastic depletion problems frequently admit a simple, myopic control policy which is at most a factor of $1/2$ to optimal. This paper presents general conditions under which a myopic control policy is a constant factor approximation algorithm for a stochastic depletion problem and goes on to verify these properties for broad families of stochastic depletion problems. This, in turn, yields myopic approximation algorithms for a number of interesting dynamic optimization applications.

There are several areas that deserve continued study. From an algorithmic perspective, one may consider k -step lookahead policies as a generalization of the myopic (1-step lookahead) policies analyzed here. Such policies select, at every point in time, an action that is optimal for a problem with a horizon precisely k time-steps ahead. It would be interesting to understand whether, or under what conditions, such policies may be expected to dominate the myopic policy.

In addition to the applications in §5 and §6, it would be interesting to explore other dynamic stochastic optimization problems that may be studied either within our framework or perhaps slight modifications to it. For instance, the generalized assignment problem (Shmoys and Tardos [26]) is known to reduce to the maximization of a submodular function over a partition matroid. An interesting stochastic generalization of this problem that would allow for a number of interesting applications would involve making the successful placement of an item in a bin stochastic. Unfortunately, this particular stochastic generalization does not reduce to a stochastic depletion problem but is nonetheless very similar to one.

Another broad issue is identifying other families of stochastic depletion problems that satisfy the VFM and IR properties, or in another direction, identifying conditions under which we may not expect one of those properties to be satisfied. Yet another issue is the optimality of our approximation schemes: for deterministic variants of several of the application problems considered in this work such as the submodular maximization problems over matroids, there exist (typically, fairly complex) *offline* algorithms that admit an approximation ratio of $e/(e - 1)$ (see Calinescu et al. [5]). This guarantee is known to be optimal. That is, no efficient approximation algorithm with a superior guarantee exists, unless $P = NP$. It would be interesting to understand whether an approximation ratio of 2 is optimal in some sense for either of the families of problems for which we have established that guarantee in this paper. In any case, given that the best approximation guarantee we may expect is a factor of $e/(e - 1)$, it is remarkable that a simple myopic scheme comes so close to achieving that guarantee and that this guarantee may be established in the generality of the stochastic depletion framework.

Acknowledgments. The authors are grateful to Retsef Levi for valuable discussions, particularly in relation to §6. Section 5.1 was motivated in part by discussions the second author had with Ben Van Roy several years ago. The second author thanks Andreas Schulz for pointing out his interesting work on submodular maximization. Finally, the authors thank two anonymous referees for useful comments that improved the exposition of this work. The first author was supported by a STMicroelectronics Stanford Graduate Fellowship. This research was supported in part by the Solomon Buchsbaum Research Fund.

Appendix. Miscellaneous technical proofs.

PROOF OF LEMMA 4.3. Consider using an optimal policy starting at state \bar{s} . Let us partition the initial set of jobs into a set of “fictitious” and “real” jobs; we assume that we begin with α_m fictitious jobs of type m and $x(\bar{s})_m - \alpha_m$ real jobs of type m . This partitioning serves purely as a labeling of jobs and does not impact the system in any fashion. In particular, if at some point in time t , we have $x_{t,m}^f$ and $x_{t,m}^r$ fictitious and real jobs of type m respectively, then using activity set A results in the completion of $X_{t,m}^f$ and $X_{t,m}^r$ fictitious and real jobs respectively, where $X_{t,m}^f$ is a Bernoulli $(x_{t,m}^f, P_{t,m}(A))$ random variable and $X_{t,m}^r$ is a Bernoulli $(x_{t,m}^r, P_{t,m}(A))$ random variable (so that $X_{t,m}^f + X_{t,m}^r \stackrel{d}{=} X_{t,m}$). The revenues earned are $\sum_m X_{t,m}^f w_{t,m}$ and $\sum_m X_{t,m}^r w_{t,m}$ and we are left with $x_{t,m}^f - X_{t,m}^f$ and $x_{t,m}^r - X_{t,m}^r$ fictitious and real jobs respectively.

Denote by $J^{f,*}(\bar{s})$ the expected reward-to-go under an optimal policy starting at state \bar{s} earned from the completion of fictitious jobs. Likewise, we define $J^{r,*}(\bar{s})$ as the expected reward-to-go under an optimal policy starting at state \bar{s} earned from the completion of real jobs. Now, by construction, $J^{f,*}(\bar{s}) + J^{r,*}(\bar{s}) = J^*(\bar{s})$. Since at best our scheduling policy can exhaust all fictitious jobs and since $w_{t,m}$ is nonincreasing in t for all m , $J^{f,*}(\bar{s}) \leq \sum_m \alpha_m w_{t(\bar{s}),m}$. Now, $J^{r,*}(\bar{s})$ may be viewed as the reward-to-go under some admissible policy π starting at state $\tilde{S}_\alpha(\bar{s})$. Noting that $x(\tilde{S}_\alpha(\bar{s}))_m$ is precisely the initial number of “real” jobs of type m , we then have: $J^{r,*}(\bar{s}) = J^\pi(\tilde{S}_\alpha(\bar{s})) \leq J^*(\tilde{S}_\alpha(\bar{s}))$. Consequently, we have

$$J^*(\bar{s}) = J^{f,*}(\bar{s}) + J^{r,*}(\bar{s}) \leq \sum_m \alpha_m w_{t(\bar{s}),m} + J^*(\tilde{S}_\alpha(\bar{s})),$$

which is the result. \square

PROOF OF LEMMA 5.1. Note that since the processes P_t^u are deterministic here, we may without loss restrict attention to policies that are functions of only time t and (x^1, x^2, \dots, x^U) . Here $x^u(s)_i$ is equal to 1 if in state s user u has not yet received page i . Let $w_i^u = r_m$ for item $m = (i, u, \tau_m, d_m)$. Recall by assumption that all packets are known a priori and share a common deadline of T so that $\tau_m = 0$ and $d_m = T$ for all m . We define the set of myopic packets as $\mathcal{P}_0^* = \arg \max_i \sum_u C_u w_i^u x^u(s_0)_i$. Let us assume for the sake of contradiction that in state s_0 no optimal policy transmits a packet in \mathcal{P}_0^* . Let π^* be an optimal policy; $\pi^*(s_0) \notin \mathcal{P}_0^*$. Without loss of generality, we can assume that $t(s_0) = 0$. In particular, if $t(s_0) \neq 0$, we may consider a problem with time indexed by $t' = t - t(s_0)$ and horizon $T' = T - t(s_0)$.

Define a policy $\bar{\pi}$ according to:

$$\begin{aligned} \bar{\pi}(s) &= \pi^g(s) \quad \forall s \quad \text{s.t.} \quad t(s) = 0, \\ \bar{\pi}(s) &= \pi^*(f(s)) \quad \forall s \quad \text{s.t.} \quad t(s) > 0, \end{aligned}$$

where $f: \mathcal{S} \rightarrow \mathcal{S}$ is defined according to $x(f(s))_{\pi^g(s_0)} = x(s_0)_{\pi^g(s_0)}$, $x(f(s))_i = x(s)_i \quad \forall i \neq \pi^g(s_0)$, and $t(f(s)) = t(s) - 1$.

Letting S_t^* denote the random state at time t under π^* , let $\tau = \min\{t < T: \pi^*(S_t^*) \in \mathcal{P}_0^*\}$ (set τ to ∞ if the set is empty). Since $P_t^u = C_u$ for all t , we may, for an arbitrary alternative policy π , couple the sample paths under π and π^* such that $Y_{i,k}^\pi = Y_{i,k}^{\pi^*}$ for all k, i , where $Y_{i,k}^\pi$ and $Y_{i,k}^{\pi^*}$ denote, respectively, the random vector of successful transmissions of packet i at the k th transmission of that packet under policy π and π^* , respectively. Consider using policy $\bar{\pi}$ for $t \leq \tau$ and policy π^* thereafter. We call this policy π' . π' may be constructed as an admissible randomized policy. Letting S'_t denote the random state in time slot t under π' , we have under this coupling that $S'_t = S_t^*$ for all $t \geq \tau + 1$. For $t < \tau + 1$, $x(S'_{t+1})_k = x(S'_t)_k$ for all $k \neq \pi^g(s_0)$. $x(S'_t)_{\pi^g(s_0)} = x(s_0)_{\pi^g(s_0)} - Y_{\pi^g(s_0),1}^{\pi'}$ for $0 < t < \tau + 1$, while $x(S'_t)_{\pi^g(s_0)} = x(s_0)_{\pi^g(s_0)}$ for $0 < t < \tau$.

Denote by $R_t^{\pi'}$ and $R_t^{\pi^*}$ the random rewards earned in the t th time-step under the π' and π^* policies respectively, so that $J^*(s_0) = E[\sum_t R_t^{\pi^*}]$ and $J^{\pi'}(s_0) = E[\sum_t R_t^{\pi'}]$. Now observe that by our construction, the following facts are true.

(i) $E[R_0^{\pi'}] \geq E[R_{\tau \wedge T-1}^{\pi^*}]$. Recall that since there are no arrivals in this system, $x(S_{\tau \wedge T-1}^*) \leq x(s_0)$ for all i . We have:

$$\begin{aligned} E[R_{\tau \wedge T-1}^{\pi^*}] &\leq E\left[\max_i E[R(S_{\tau \wedge T-1}^*, i) \mid S_{\tau \wedge T-1}^*]\right] \\ &= E\left[\max_i \sum_u C_u w_i^u x^u(S_{\tau \wedge T-1}^*)_i\right] \\ &\leq E\left[\max_i \sum_u C_u w_i^u x^u(s_0)_i\right] \\ &= \max_i \sum_u C_u w_i^u x^u(s_0)_i \\ &= \max_i E[R(s_0, i)] \\ &= E[R_0^{\pi'}], \end{aligned}$$

where the first equality follows from the fact that $\tau \wedge T - 1$ is a stopping time w.r.t. $\sigma(S_0^*, \dots, S_t^*)$.

(ii) $R_t^{\pi'} = R_{t-1}^{\pi^*}$ for $0 < t \leq \tau$ by the definition of π' which uses policy $\bar{\pi}$ for $0 < t \leq \tau$ and by our coupling.

(iii) $R_t^{\pi'} = R_t^{\pi^*}$ for $t > \tau$ because by our coupling and the definition of π' , $S_{\tau+1}^{\pi'} = S_{\tau+1}^{\pi^*}$. Then, for $t > \tau$, π' and π^* coincide.

We thus have:

$$\begin{aligned} \sum_{t=0}^{T-1} R_t^{\pi'} &= R_0^{\pi'} + \sum_{t=1}^{T-1} [1_{t \leq \tau} R_{t-1}^{\pi^*} + 1_{t > \tau} R_t^{\pi^*}] \\ &= R_0^{\pi'} + \sum_{t=1}^{T-1} [1_{t \leq \tau} (R_{t-1}^{\pi^*} - R_t^{\pi^*}) + R_t^{\pi^*}] \\ &= R_0^{\pi'} - 1_{T-1 < \tau} R_{T-1}^{\pi^*} - 1_{T-1 \geq \tau} R_{\tau}^{\pi^*} + \sum_{t=0}^{T-1} R_t^{\pi^*} \\ &= R_0^{\pi'} - R_{T-1 \wedge \tau}^{\pi^*} + \sum_{t=0}^{T-1} R_t^{\pi^*}, \end{aligned}$$

where the first equality follows from points ii and iii. Taking expectations yields

$$\begin{aligned} J^{\pi'}(s_0) &= E\left[\sum_{t=0}^{T-1} R_t^{\pi'}\right] = E[R_0^{\pi'}] - E[R_{T-1 \wedge \tau}^{\pi^*}] + E\left[\sum_{t=0}^{T-1} R_t^{\pi^*}\right] \\ &\geq E\left[\sum_{t=0}^{T-1} R_t^{\pi^*}\right] \\ &= J^*(s_0), \end{aligned}$$

where the inequality follows from point i. Thus, π' is an optimal policy as well. This contradicts our assumption that no optimal policy transmits a packet in \mathcal{P}_0^* in state s_0 . We may thus assume without loss that for all states $s_0 \in \mathcal{S}$ an optimal policy transmits a packet in \mathcal{P}_0^* . This suffices for the proof. \square

References

- [1] Bar-Noy, A., S. Guha, Y. Katz, J. Naor, B. Schieber, H. Shachnai. 2002. Throughput maximization of real-time scheduling with batching. *Proc. 13th SODA*, San Francisco. 742–751.
- [2] Bassamboo, A., J. M. Harrison, A. Zeevi. 2006. Design and control of a large call center: Asymptotic analysis of an LP-based method. *Oper. Res.* **54** 419–435.
- [3] Bassamboo, A., J. M. Harrison, A. Zeevi. 2006. Dynamic routing and admission control in high volume service systems: Asymptotic analysis via multi-scale fluid limits. *Queueing Systems Theory Appl.* **51** 249–285.
- [4] Buyukkoc, C., P. Varaiya, J. Walrand. 1985. The $c - \mu$ rule revisited. *Adv. Appl. Probab.* **17**(1) 237–238.
- [5] Calinescu, G., G. Chekuri, M. Pál, J. Vondrák. 2007. Maximizing a submodular set function subject to a matroid constraint. *Proc. 12th IPCO*, Ithaca, NY, 182–196.
- [6] Chou, P. A., Z. Miao. 2006. Rate-distortion optimized streaming of packetized media. *IEEE Trans. Multimedia* **8**(2) 390–404.
- [7] Dua, A., N. Bambos. 2006. Downlink scheduling of heterogeneous traffic. *Proc. IEEE ICC*, Istanbul, Turkey, Vol. 11, 5252–5257.

- [8] Dua, A., N. Bambos. 2006. Joint power allocation and scheduling for deadline constrained wireless traffic. *Proc. IEEE Globecom* November 2006, 1–5.
- [9] Dua, A., N. Bambos. 2007. Downlink wireless packet scheduling with deadlines. *IEEE Trans. Mobile Comput.* **6**(12) 1410–1425.
- [10] Eryilmaz, A., R. Srikant, J. R. Perkins. 2005. Stable scheduling policies for fading wireless channels. *IEEE/ACM Trans. Netw.* **13**(2) 411–424.
- [11] Fisher, M. L., G. L. Nemhauser, L. A. Wolsey. 1978. An analysis of approximations for maximizing submodular set functions—II. *Math. Programming Study* **8** 73–87.
- [12] Fleischer, L. K., M. X. Goemans, V. Mirrokni, M. Sviridenko. 2006. Tight approximation algorithms for maximizing general assignment problems. *Proc. 17th SODA*, Miami, 611–620.
- [13] Frank, R. E., W. F. Massy, Y. Wind. 1972. *Market Segmentation*. Prentice-Hall, Englewood Cliffs, NJ.
- [14] Gandhi, T., S. Khuller, S. Parthasarathy, A. Srinivasan. 2002. Dependent rounding in bipartite graphs. *Proc. 43rd FOCS*, Vancouver, B.C., Canada, 323–332.
- [15] Goundan, P. R., A. S. Schulz. 2007. Revisiting the greedy approach to submodular set function maximization. Working paper, Massachusetts Institute of Technology, Cambridge, MA.
- [16] Harrison, J. M., A. Zeevi. 2005. A method for staffing large call centers using stochastic fluid models. *Manufacturing Service Oper. Management* **7** 20–36.
- [17] Hopp, W. J., X. Xu. 2005. Product line selection and pricing with modularity. *Manufacturing Service Oper. Management* **7**(3) 172–187.
- [18] Huang, J., R. Berry, M. Honig. 2005. Wireless scheduling with hybrid ARQs. *IEEE Trans. Wireless Comm.* **4**(November 6) 2801–2810.
- [19] Jansen, K., G. Zhang. 2004. On rectangle packing: Maximizing benefits. *Proc. 17th SODA*, New Orleans, 204–213.
- [20] Kim, J.-H., K.-Y. Chwa. 2004. Scheduling broadcasts with deadlines. *Theoretical Comput. Sci.* **325**(3) 479–488.
- [21] Kohli, R., R. Sukumar. 1990. Heuristics for product-line design using conjoint analysis. *Management Sci.* **36**(12) 1464–1478.
- [22] Moorthy, K. S. 1984. Market segmentation, self-selection, and product line design. *Marketing Sci.* **3**(4) 288–307.
- [23] Nemhauser, G. L., L. A. Wolsey, M. L. Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Math. Programming* **14** 265–294.
- [24] Pigou, A. C. 1978. *The Economics of Welfare*. AMS Press, Inc., NY.
- [25] Ren, T., I. Koutsopoulos, L. Tassiulas. 2002. QoS provisioning for real-time traffic in wireless packet networks. *Proc. IEEE GLOBECOM* **2**(November) 1673–1677.
- [26] Shmoys, D. B., E. Tardos. 1993. An approximation algorithm for the generalized assignment problem. *Math. Programming* **62**(3) 461–474.
- [27] Su, C.-J., L. Tassiulas. 1997. Broadcast scheduling for information distribution. *Proc. IEEE INFOCOM* **1**(April) 109–117.
- [28] Tsibonis, V., L. Geogiadis. 2005. Optimal downlink scheduling policies for slotted wireless time-varying channels. *IEEE Trans. Wireless Comm.* **4**(4) 1808–1817.
- [29] van Mieghem, J. A. 1995. Dynamic scheduling with convex delay costs: The generalized $c | \mu$ rule. *Ann. Appl. Probab.* **5**(3) 809–833.
- [30] van Ryzin, G. J., S. Mahajan. 1999. On the relationship between inventory costs and variety benefits in retail assortments. *Management Sci.* **45**(11) 1496–1509.
- [31] Yunes, T. H., D. Napolitano, A. Scheller-Wolf, S. Tayur. 2007. Building efficient product portfolios at John Deere and Company. *Oper. Res.* **55**(4) 615–629.