

Model Predictive Control for Dynamic Resource Allocation

Dragos Florin Ciocan, Vivek Farias

Massachusetts Institute of Technology Sloan School of Management, Cambridge, Massachusetts 02139

{ciocan@mit.edu, vivekf@mit.edu}

The present paper develops a simple, easy to interpret algorithm for a large class of dynamic allocation problems with unknown, volatile demand. Potential applications include ad display problems and network revenue management problems. The algorithm operates in an online fashion and relies on reoptimization and forecast updates. The algorithm is robust (as witnessed by uniform worst-case guarantees for arbitrarily volatile demand) and in the event that demand volatility (or equivalently deviations in realized demand from forecasts) is not large, the method is simultaneously optimal. Computational experiments, including experiments with data from real-world problem instances, demonstrate the practicality and value of the approach. From a theoretical perspective, we introduce a new device—a balancing property—that allows us to understand the impact of changing bases in our scheme.

Key words: model predictive control; stochastic control; approximation algorithms; online advertising

MSC2000 subject classification: Primary: 93E20, 90C39, 90C59; secondary: 91B32, 90B22, 90B60

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

History: Received May 4, 2011; revised September 11, 2011. Published online in *Articles in Advance* June 19, 2012.

1. Introduction. The present paper considers an archetypal dynamic allocation problem that captures a swathe of disparate applications in revenue management and e-commerce. Informally, the class of problems we consider can be described as follows: we are given a bipartite graph consisting of a set of I sources and A sinks, together with E edges from the sources to the sinks. Every source node i receives “demand” over time. The rate at which demand arrives is described by a general stochastic process. An allocation is an assignment of the demand arriving at each of the source nodes to the sinks they are connected with. This allocation may change over time; hence the dynamic moniker. For every unit of demand allocated along edge e (connecting some source $i(e)$ to some sink $a(e)$), we receive a reward of p_e . In addition, this unit allocation will consume varying quantities of each of K distinct resources described by a vector $A_e \in \mathbb{R}_+^K$. We begin with an initial allocation of each of the K resources and must allocate demand over time in a manner that maximizes revenues while consuming no more than the initial allocation of resources. The key source of uncertainty in this model comes from the stochastic demand rate processes; in practical settings even specifying such a process is a potentially nontrivial task.

The abstract allocation model we describe can capture a number of applications of broad interest. Two examples that are of particular interest in revenue management (RM) are the following:

(i) *The network revenue management (NRM) problem.* This is a generic high-dimensional allocation problem encountered in industries ranging from the hospitality industry to the airline industry and is effectively a cornerstone RM model. For concreteness, consider the problem faced by an airline that sells a variety of itineraries over a network of cities over time. Each itinerary requires seats on different legs of the network and generates different revenues. The airline must sell itineraries over time in a manner that respects seat capacity and maximizes revenues. Volatility in demand and the high-dimensional nature of the allocation problem together make it challenging.

(ii) *Online ad display problems.* An ad network serves as an intermediary between publishers (sources of traffic or demand) and advertisers. Via a variety of contractual agreements it agrees to display ads from specific advertisers to compatible traffic from publishers (where the notion of compatible might correspond to the publisher’s entity among other things). Advertising contracts might specify that a given ad is displayed up to a certain number of times to compatible traffic. Alternatively, by specifying a dollar rate for the display of a specific ad to a specific type of demand, the advertiser might commit to spending up to a certain budget on advertising. The ad exchange may, in turn, seek to exhaust as much of this budget as possible while maximizing revenues. Again, volatility in traffic across publishers and over time, combined with the sheer number of potential alternatives for the allocation of a given unit of traffic, makes this a nontrivial allocation problem.

We defer a rigorous problem definition and a precise explanation of how these problems and others fit into our framework to §2. For now, we simply note that the above problems are considered difficult primarily because of uncertainty in demand and the high-dimensional nature of the resource allocation problem at hand. As such, there are distinct bodies of literature devoted to the problems and variants thereof. Rigorous algorithmic approaches appear to fall roughly into two categories at diametrically opposite ends of the spectrum on demand assumptions. On the one hand, by assuming that the demand processes admit only “small” shocks so that uncertainty in total

demand at each source is small relative to its magnitude,¹ one may resort to solving simple “offline” versions of the allocation problem at hand. These solutions yield, via an appeal to a law of large numbers argument, an essentially optimal solution to the actual allocation problem. This has been the dominant modeling approach in the majority of RM applications. On the other hand, the assumption of small demand shocks is clearly an idealization, so that a distinct approach to such problems has been to assume that demand is adversarial. One then seeks to design online allocation schemes that compete effectively with the adversary generating demand. This has been the dominant modeling approach for many of the e-commerce related allocation problems we allude to. Although such online schemes are typically quite simple, their design and analysis is typically fairly brittle to model assumptions.

In reality one typically faces a world that is somewhere in between these assumptions: whereas assuming that demand is effectively deterministic (by assuming small shocks) is potentially unrealistic, the assumption of adversarial demand is itself potentially conservative and unrealistic. In particular, in many instances of the applications discussed, one generates copious amounts of historical data on demand over time that in principle might be used to construct useful forecast models. Unfortunately, such forecast models are far from easy to construct in practice—challenges include judging the historical relevance of data, learning factors that serve as useful predictors, and of course, problem scale. Moreover, even assuming access to such a forecast model, the resulting dynamic optimization problem remains high dimensional and intractable. The present paper tackles this middle ground on demand assumptions and provides solutions that attempt to address these challenges.

1.1. Contributions. The present work posits a simple new approach to the general class of allocation problems previously described that relies on a combination of reoptimization and “robust” forecasting; importantly, the approach does not require that the demand rate process be specified. The approach is pragmatic and at the same time admits attractive theoretical performance guarantees. In greater detail, we make the following contributions:

(i) *A simple allocation scheme.* We design an allocation scheme that relies on frequent reoptimization using suitably updated forecasts. In particular, at discrete points in time, one uses demand realized up to that point in time to construct a forecast for future demand in a precisely specified fashion. Assuming these forecasts to be exact, one then solves a simple linear optimization problem that prescribes an allocation of demand from sources to sinks. This allocation rule is followed until the next opportunity to reoptimize. The scheme is entirely mechanical in that it can be applied to any dynamic allocation problem within our framework without any instance specific analysis.

(ii) *Uniform worst-case guarantees.* Assuming that the demand rate process faced lies in a certain broad class of stochastic processes,² we show that our allocation scheme yields expected revenues that are within a constant factor of expected revenues under a certain super-optimal policy. The value of this constant is either 0.342 or 0.2 depending on the specific assumptions we make on the family of demand processes. These worst-case results are remarkable—they hold for *arbitrarily volatile* demand processes and illustrate that the proposed scheme is robust across a broad class of demand processes while being oblivious to the specification of the process. Our performance analysis overcomes the technical hurdle of analyzing the impact of basis changes in certain math programs that underlie our allocation scheme; this is the primary hurdle in analyzing allocation schemes that rely on reoptimization.

(iii) *Parametric guarantees.* In addition to uniform worst-case guarantees, we present performance guarantees that reveal that as the volatility of the underlying demand process shrinks, our allocation scheme approaches optimality. Together with the worst-case guarantees, this allows the following interpretation: the scheme we propose is essentially optimal if available forecasts are accurate but otherwise robust to forecast inaccuracies.

(iv) *Computational evidence.* We present computational experiments on both synthetic problem instances as well as a real-world example of an ad display problem using demand data from a mobile ad network. In both the synthetic and real-world instances, the proposed scheme is seen to provide performance levels that are typically well within 90% of an upper bound constructed by assuming that demand realizations were available a priori.

Although our literature review will extensively place these contributions in the context of the extant work on dynamic allocation problems, we end this discussion with a brief statement of the relative merits: The nature of the assumptions made in traditional RM modeling (that of small demand shocks) translates in our model to

¹ This is analogous to assuming a deterministic demand rate process in the general model we will consider.

² We essentially allow for multivariate Gaussian processes with continuous sample paths. We require that the volatility of any variate be a concave increasing function of time; we show that these requirements may be viewed as natural in the context of any stochastic forecast model.

a demand rate process that is deterministic; the assumption of a *stochastic* demand process with no restrictions on volatility allows one to model (and address) large shocks in demand. Conversely, if one were to adopt an adversarial view of demand within our model, the nature of an optimal online scheme varies considerably across specializations of our model as do the corresponding competitive ratios. Loosely speaking, optimal competitive ratios range from a constant (in certain variants of the ad display problem) to scaling inversely with the log of the number of demand types (in the case of a general packing problem such as the NRM problem); the corresponding allocation schemes are quite distinct. Moreover, one typically loses the notion of what it means to have an “accurate” forecast model in such a setting. In contrast, we present a unified approach based on reoptimization and forecast updates to what is apparently a broad range of problems and establish that it performs well in an expected sense assuming a broad generative family of demand processes.

1.2. Literature review. By virtue of its generality, our dynamic allocation model bears comparison to a very broad class of models. As such we organize our review of relevant literature around generic algorithmic approaches to comparable models. Given the diversity of comparable dynamic allocation models, this review is by necessity incomplete and biased.

Online algorithms. The dynamic allocation problems we study are quite similar in spirit to (multidimensional) online packing problems. A powerful tool in the design of schemes for such problems is the primal-dual schema. Under an assumption of entirely adversarial demand it is possible to use this schema to design online algorithms that bear a competitive ratio on the order of the logarithm of the number of item types (i.e., $O(\log I)$ in the context of our model); see Buchbinder and Naor [10]. Under the so-called random permutation model, where an adversary selects the number of arriving items but nature then permutes these items uniformly at random, substantially stronger guarantees are possible—in particular, it is possible to provide online algorithms that constitute polynomial time approximation schemes under this model for a variety of packing problems in appropriate regimes. In particular, this was pointed out by Kleinberg [18] in the context of a secretary problem. More recently, Agrawal et al. [2] developed an online polynomial-time approximation scheme for a general multidimensional packing problem under this model. The exchangeable nature of the demand distribution under this adversarial model is what effectively drives these results; unfortunately this sort of model does not appear appropriate for demand processes that are inherently nonstationary.

There has also been a good amount of work on online algorithms for some of the specific applications we consider. In particular, for the so-called AdWords problems (a specialization of the ad display problem we discuss), Mehta et al. [21], design an online $(1 - 1/e)$ -competitive algorithm. Mahdian et al. [20] work with the same Adwords problem, but allow for the existence of forecasts of demand (keyword arrivals). They not only obtain a constant factor guarantee versus worst-case (adversarial) inputs but also a constant factor guarantee versus the revenues that would be achieved if the forecast was perfect. The analysis in both papers assumes individual bids are small compared to the total budgets. In the case of the NRM problem, Ball and Queyranne [4] consider a simplified version of the NRM problem, namely, a setting in which the airline operates multiple itineraries on a *single* leg, i.e., $K = 1$ and show that a competitive ratio of two is achievable and optimal.

Many of the online schemes we have discussed are inherently conservative, which leads to hesitation in their adoption in practice.

Approximate dynamic programming (ADP). Given a model of demand, one may in principle solve the sort of dynamic allocation problems we study via dynamic programming. Of course, this is untenable in practice because of the curse of dimensionality, and one approach of contending with this issue is the design of appropriate ADP schemes. For instance, Bertsimas and Demir [7] solve integer multidimensional knapsack problems via ADP. ADP heuristics have also been developed for the NRM problem. In particular, see Adelman [1], Farias and Van Roy [13], and Zhang and Adelman [24]. Although with careful tailoring to the problem in question, these heuristics often exhibit excellent practical performance, (absolute) theoretical performance guarantees are difficult to come by in general. Moreover, the “tailoring” required is frequently nontrivial—one needs to provide these algorithms with good “approximation architectures” for the problem at hand.

Fluid models. Fluid models arise essentially from considering allocation problems with their (stochastic) demand processes replaced by “fluid” arrival processes with deterministic rates matching those of the stochastic demand process; of course, these rates have to be available a priori. Solving these models is typically substantially easier than the original stochastic problem and the resulting solutions can work well in the original stochastic problem. Briefly, one can expect good performance provided the deviation in cumulative demand relative to its mean is small, and this typically requires a suitable scaling of the original problem. The model we study is, in contrast, a *stochastic* fluid model where the rates of the fluids in question are stochastic processes as opposed

to being deterministic and known a priori. Said another way, our work can be viewed as taking the fluid model approach *without* a priori information on the demand rate process, which is itself stochastic.

A typical area of application is in the analysis of control policies for queuing networks; see for instance Bramson [9]. A second area that has found applications for these tools is revenue management. For instance, the seminal work of Gallego and van Ryzin [15] essentially posits a fluid model for the NRM problem and then proceeds to show that solutions derived from this model work well assuming the scale of demand and capacity grow large simultaneously. It is worth reiterating that this again requires one know the *rate* of the underlying demand process a priori and in that sense such a model is unable to capture “large” shocks in demand. There is a vast literature preceding this paper, two of which do make an attempt to capture large shocks in demand—the first is a paper by Akan and Ata [3] that considers a stochastic fluid model for NRM. Their model is very closely related to the one studied here. The authors provide a remarkable characterization of optimal policies and show how to compute these optimal policies assuming the demand process is described by a diffusion. Unfortunately, this still requires that one specify the diffusions a priori. A second paper by Chen and Farias [12], which is perhaps most closely related to the present work, studies a *one* dimensional allocation problem (with a somewhat distinct control/reward mechanism) under similar assumptions as those studied here. We present performance guarantees relative to a multivariate version of the family of demand processes studied in that work.

Reoptimization/model predictive control. The area of model predictive control essentially prescribes solving hard stochastic control problems by posing fluid models and then *resolving* these fluid models with suitable updates as uncertainty reveals itself. Our approach falls squarely within this philosophy. For a survey of model predictive control literature, see Bemporad [5].

The MPC approach has been used (under different names) in a variety of settings. One prime example is in scheduling for queuing networks; for instance, the seminal work of Chen and Yao [11] can be viewed in this light as can the celebrated max weight scheduling policy (see Shah and Wischik [23]). More recently, this philosophy has also found application in revenue management, and in particular, for the NRM problem. In that context, Maglaras and Meissner [19] pointed out that repeatedly re-solving the fluid model linear program prescribed by Gallego and van Ryzin [15] yields an optimal allocation in the fluid scale. Subsequent work has shown that this reoptimization can play a substantial role in accelerating the rate at which such allocations approach optimality as the problem is scaled. In particular, Reiman and Wang [22] show that a single reoptimization at a carefully chosen point can result in an additive performance loss that grows like the square root of the problem scale. Recently, Jasin and Kumar [16] provided an extremely elegant demonstration of the fact that with repeated reoptimization (that can be at uniformly spaced intervals) the additive performance loss is *independent* of the problem scale. This work nicely illustrates the impact of reoptimization in combating small shocks in demand (the demand rate is *known* in all of the previously mentioned papers). The present paper can be seen as complementing that work by showing that reoptimization *with* appropriate forecast updates is beneficial when the demand rate itself is unknown and stochastic and thereby aids in combating large shocks.

2. Model. This section will describe our model rigorously. We then present examples of three problems that are captured within our model. The first two are the NRM and ad display problems described loosely earlier. The third concerns revenue management in a multiclass processing network with stochastic arrival rates.

System. Consider a bipartite graph with I sources indexed by i and S sinks indexed by a . The edge set of this graph has size E , and a generic edge will be denoted by e . We understand by $i(e)$ the source node for edge e and by $a(e)$ its sink node. Given this graph, which will underlie a general allocation process over time from sources to sinks, we next describe several key model primitives:

(i) *Demand.* We associate each source with a nonnegative real-valued stochastic process, $\{\Lambda_{i,t}(\omega)\}$, with continuous sample paths.³ Each of these I processes are defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The total demand at source i over the horizon $[0, T]$ will be understood as $\int_0^T \Lambda_{i,t} dt$. We denote by \mathcal{F}_t , the sigma algebra generated by the sample paths of the I demand processes up to time t , i.e., $\mathcal{F}_t = \sigma(\{\Lambda_s(\omega) : 0 \leq s \leq \hat{t}, \hat{t} \leq t, \omega \in \Omega\})$.

(ii) *Resources and resource consumption.* We are given a set of K distinct resources indexed by k . The available quantity of these resources at time t is given by a vector $x_t \in \mathbb{R}_+^K$. A unit allocation of demand along edge e will consume resources. The amount of resource k consumed is given by $A_{k,e}$; the vector of resources consumed by a unit allocation on edge e is thus the column vector $A_{\cdot,e} \triangleq A_e$. Denote by $A \in \mathbb{R}_+^{K \times E}$ the matrix whose e th column is A_e .

³ We will suppress the dependence on ω if this is clear from context.

(iii) *Prices/revenues.* Allocating a unit of demand along edge e generates revenue p_e . Denote by $p \in \mathbb{R}_+^E$ the column vector whose e th component is p_e .

Control. Although our system evolves in continuous time, control is exerted at discrete times, $\{0, T/N, 2T/N, \dots, T\} \triangleq \mathcal{T}_N$. The control chosen at time iT/N remains in effect over the interval $[iT/N, (i+1)T/N)$ and is specified by a vector $z \in \mathcal{Z}$ where the set of feasible controls at time t , \mathcal{Z} , is defined as

$$\mathcal{Z} = \left\{ z \in \mathbb{R}_+^E : \sum_{e: i(e)=i} z_e \leq 1 \ \forall i \right\}.$$

The control chosen at time iT/N determines the allocation of demand across edges over the subsequent interval. An *admissible control policy* is a \mathcal{Z} -valued process adapted to the filtration \mathcal{F}_t , which we will denote $\{z_t\}$. Denote the set of admissible control policies by Π_N . We also denote by $\{z_{e,t}\}$ the e th component of $\{z_t\}$, i.e., the process describing the allocations across a particular edge e .

Dynamics. The state of the system at time t is effectively the history of the exogenously evolving demand processes up to that point and the quantities x_t of the K resources that remain. The evolution of x_t is specified by the differential equation:

$$\frac{d}{dt} x_{k,t} = \begin{cases} -\sum_e A_{k,e} \Lambda_{i(e),t} z_{e,d(t)} & \text{if } x_{k,t} > 0 \\ 0 & \text{otherwise} \end{cases}$$

for all k . Here $d(t) = \max_{\{i: iT/N \leq t\}} iT/N$.

The problem. Define the event $I_{e,t} = \{x_{k,t} > 0 \ \forall k \text{ s.t. } A_{k,e} > 0\}$. We are tasked with finding an admissible control policy $\{z_t\}$ that solves the following optimization problem:

$$\max_{\{z_t\} \in \Pi_N} \mathbb{E} \left[\int_0^T \sum_e p_e \Lambda_{i(e),t} z_{e,d(t)} \mathbf{1}_{\{I_{e,t}\}} dt \right]. \tag{1}$$

We denote the optimal value to this optimal control problem by $J^{*,N}(x_0)$.⁴

The model we consider is best thought of as a *stochastic fluid model*. In particular, in a real system demand is potentially best captured as a multivariate point process with rate Λ_t ; in our model we ignore the fluctuations of this counting process from its mean focusing instead on shocks in the rate process itself. In most applications the fluctuations of the counting process about its mean are small in that their effect can be shown to be negligible in a regime where x_0 and Λ_t are scaled simultaneously by some scale factor that grows large.

2.1. A family of stochastic demand models. Although the algorithms we present will apply to general nonnegative rate processes, it is hopeless to expect any single algorithm to perform well across a class this broad. As such our *analysis* will be limited to studying the performance of our prescriptions across a more limited family of processes. We consider this family here and then present examples of processes within this family:

ASSUMPTION 2.1. *Structure of $\{\Lambda_t\}$:*

- (i) $\Lambda_{i,t} = (\bar{\Lambda}_{i,t})^+$, where $\bar{\Lambda}_t$ is a Gaussian process with continuous sample paths.
- (ii) $\mathbb{E}[\Lambda_{i,t}] \triangleq \lambda_{i,t}$ is positive.
- (iii) The variance of the random variable $\bar{\Lambda}_{i,t}$, $\sigma_{i,t}^2$, is nondecreasing as a function of t and concave.

Note that we do not make any assumptions on the correlation structure for this multivariate stochastic process. Although restricting attention to rate processes within some class of stochastic processes is certainly restrictive relative to, say, an adversarial model of rate processes, the family of processes permitted by Assumption 2.1 is ubiquitous in applications.

2.1.1. Ubiquity of assumption 2.1. We now demonstrate a family of processes that is both ubiquitous in applications and satisfies the requirements of the Assumption 2.1. Consider processes $\{\bar{\Lambda}_t\}$ defined according to

$$\bar{\Lambda}_{i,t} = \lambda_{i,t} + \int_0^t \phi_i(t-s) dZ_{i,s},$$

where $\lambda_{i,t}$ is some Lipschitz-continuous, nonnegative function of t for all i , $\phi_i(\cdot)$ is a Lipschitz-continuous function of t such that $|\phi_i(\cdot)|$ is nonincreasing, and Z_s is I -dimensional Brownian motion. For reasons that will become apparent shortly, we will refer to these as moving average processes. Moving average processes satisfy the requirements of Assumption 2.1.

⁴ We will often omit a reference to N when it is clear from context.

To begin, we note that the class of moving average processes include a fair number of common continuous time processes used in stochastic modeling. For instance, setting $\phi_i = 1$ yields the Wiener process, and setting $\phi(s) = \exp(-s)$ yields the well-known Ornstein-Uhlenbeck (or Langevin) process. The ubiquity of this class is perhaps more obvious if one considers evaluating the process at discrete points in time.

In particular, for any $\Delta > 0$, we have (via Ito isometry) that

$$\bar{\Lambda}_{i, n\Delta} = \lambda_{i, n\Delta} + \sum_{k=0}^{n-1} \theta_{n-k} \epsilon_k,$$

where ϵ_k are independent standard normal random variables, and $\theta_j = \sqrt{\int_{(j-1)\Delta}^{j\Delta} \phi^2(s) ds}$. This is nothing but a moving average model that finds wide application in time-series modeling. In particular, Assumption 2.1 can thus be seen to permit moving average models that are

- (i) of arbitrary order;
- (ii) not necessarily stationary;
- (iii) satisfying the property that the weights $\{\theta_i\}$ be nonincreasing so that past shocks have smaller influence on the process than more recent shocks do.

This is evidently a broad class of models.

In subsequent performance analyses we will heavily exploit the natural *symmetry* in the marginals of this class of processes, as is made precise by the following lemma. The fact that it is the symmetry of the marginals of the demand process that drive our performance results is something of a unique insight.⁵

LEMMA 2.1. *Let $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be continuous, nondecreasing, and concave with $f(0) = 0$. Then, provided the process $\{\Lambda_t\}$ satisfies Assumption 2.1 and further, $\lambda_{i,t} = \lambda_i, \forall i \in \mathcal{I}, t \in [0, T]$, we must have*

$$\frac{\mathbb{E}[(1/T) \int_0^T f(\Lambda_{i,t}) dt]}{f(\mathbb{E}[(1/T) \int_0^T \Lambda_{i,t} dt])} \geq 0.342$$

for all i .

REMARK 2.1 (LOW VOLATILITY). The previous result demonstrates an outcome of the symmetry of the marginals of our stochastic process and the resultant uniform bound holds irrespective of the magnitude of σ_t . Of course, one may hope that, if σ_t is small, then one might expect a tighter bound. In fact, Lemma 7 of Chen and Farias [12] establishes that if $\sigma_{i,t}/\lambda_i \leq \sqrt{2\pi}B$ for all t , then one has

$$\frac{\mathbb{E}[(1/T) \int_0^T f(\Lambda_{i,t}) dt]}{f(\mathbb{E}[(1/T) \int_0^T \Lambda_{i,t} dt])} \geq \frac{1}{1+B} - \frac{B}{1+B} (\exp(1/4\pi B^2) + 0.853).$$

The first allocation scheme we consider *will not require* any knowledge of the specification of the process Λ_t . The second scheme we consider will use information on the drift of the process $\bar{\Lambda}_t$; in particular, this scheme will know $\{\lambda_{i,t}\}$ for all i . Note that even in the latter case, the information utilized is mild—it is trivial to construct processes with identical $\{\lambda_{i,t}\}$ but drastically different behavior; for instance, an Ornstein-Uhlenbeck process and a Wiener process.

2.2. Examples. The generic allocation problem is quite rich and encapsulates several important classes of problems described next. The first two classes collectively drive billions of dollars in commerce.

Network revenue management. Here each of the source nodes corresponds to an arriving customer class so that $\Lambda_{i,t}$ captures the arrival rate of that class. Each of the sink nodes corresponds either to an itinerary or an “offer set” depending on the precise NRM model we wish to capture. Every source has an edge to every sink. K corresponds to the number of legs on the network so that x_k is the residual capacity on leg k . We consider two distinct NRM models and interpret the control $\{z_t\}$ in each:

- (i) *Separable demand.* Here each customer class (i.e., source node) is connected to exactly one itinerary (i.e., sink). $A_{k,e}$ corresponds to the number of seats consumed on leg k by itinerary $a(e)$ and p_e is the price of this itinerary. The demand at a given source node is consequently interpreted as demand for a given itinerary-price pair. The control $\{z_t\}$ is interpreted as follows: $z_{t,e}$ is the probability that the itinerary-price pair corresponding to edge e is made available at time t . This is the predominant model for NRM problems.

⁵ In fact, analogues of Lemma 2.1 for non-Gaussian marginals drive analogous performance guarantees to those we shall derive in subsequent sections and we will provide one such example in §3.5.1.

(ii) *Customer choice.* Here a given customer class may be connected to multiple sinks. A given sink is associated with an offer set, i.e., a set of itineraries the customer might be offered from which she will pick an option desirable to her. The quantity $A_{k,e}$ corresponds to the expected number of seats consumed on leg k when customers of class $i(e)$ are presented the offer set $a(e)$; p_e is the expected price of the itinerary selected from the offer set. The control $\{z_t\}$ is then interpreted as follows: $z_{t,e}$ is the probability that the offer set corresponding to edge e is shown to an arriving customer of class $i(e)$ at time t .

What is of note in our model is the ability to capture the fact that the demand rate for a given customer class $\Lambda_{i,t}$ is *stochastic* as opposed to being deterministic (as is assumed typically). This allows us to capture “large” shocks in demand. Anticipating our scheme for allocation we will be able to do this without explicitly modeling the underlying demand process.

Online ad display allocation. Here each of the source nodes corresponds to an arriving impression type,⁶ and each of the sink nodes corresponds to an advertiser. Every source has an edge to every sink. Let K be the number of sinks; x_k can have several interpretations depending on the nature of the contract with the advertiser. For instance, the contract may specify that the advertiser will pay no more than a certain budget in which case x_k is the amount of this budget that remains. Alternatively, the advertiser may agree to having its ad displayed no more than a certain number of times over some period in which case x_k corresponds to the number of times advertiser k 's ad can be displayed. The allocation of an impression of type $i(e)$ to the advertiser $a(e)$ garners revenues p_e . In the context of budget-based contracts $A_{k,e} = p_e$; in the case of contracts based on the number of impressions served with an ad, $A_{k,e} = 1$ if $i(e)$ is an acceptable ad to advertiser $a(e)$ and zero otherwise.

It is interesting to note that past work on this problem in the adversarial setting also effectively considers a fluid model by assuming that the unit of allocation is small relative to the budget, i.e., $A_{k,e} \ll x_k$.

Revenue management for a multiclass processing network. Consider a fluid processing network wherein I types of fluids are processed. Fluid i is processed at precisely one of K processing stations at rate $1/\mu_i$. Upon being processed, a unit of fluid i is transformed into $P_{i,i'}$ units of fluid i' . Let $P \in \mathbb{R}_+^{I \times I}$ be the matrix whose i, i' th entry is $P_{i,i'}$ and assume that $I - P$ is invertible. Fluid i arrives to the system at a rate given by the stochastic process $\{\Lambda_{i,t}\}$ and the revenue from processing a unit of arriving fluid of type i (assuming all fluid generated in subsequent processing steps is also processed) is given by p_i . The goal is to design an admission policy that maximizes revenues from processing fluid over some finite horizon.⁷ This may be cast within our framework as follows: we consider a problem with I sources corresponding to each fluid type and a single sink node. For every edge e we have $p_e = p_{i(e)}$. We set x_t to be the vector of uncommitted processing time available over the remainder of the horizon at each of the K stations at time t ; thus $x_0 = T\mathbf{1}$. We set $A_e = \mu \odot (I - P)^{-1}u_{i(e)}$, where u_j is the j th unit vector.⁸ Note that the quantity $v = (I - P)^{-1}u_j$ solves the Poisson equation $v = u_j + Pv$ and represents the effective amount of fluid (of all types) that an inflow of one unit of fluid j will introduce into the system. In this setting we interpret $z_{t,e}$ as the fraction of (nonreentrant) fluid of type $i(e)$ entering the system at time t that is admitted.

It is interesting to note how this processing model departs from typical processing network models: First, the arrival rate of a fluid is stochastic; the process driving this rate is allowed to be fairly general (as discussed in the preceding section). Second, we *maximize rewards* associated with processing fluid as opposed to minimizing some cost associated with backlog. Finally, we note that we do not allow a backlog at time T . A richer formulation would allow for us to optimize some combination of rewards associated with processing fluid and costs associated with backlog at time T ; unfortunately our model does not allow this.

3. A reoptimization-based heuristic. Imagine our information structure were such that $\{\Lambda_t, t \geq 0\} \in \mathcal{F}_0$. If this were the case, the control problem at hand reduces to a deterministic optimization problem; in particular, one simply employs an allocation rule given by any optimal solution to the program

$$\begin{aligned} & \max \sum_e p_e z_e \int_0^T \Lambda_{i(e),t} dt \\ & \text{subject to } \sum_e A_{k,e} z_e \int_0^T \Lambda_{i(e),t} dt \leq x_{k,0} \quad \forall k, \\ & z \in \mathcal{Z}. \end{aligned} \tag{2}$$

⁶ An impression can be thought of as a unit of Web traffic satisfying certain preassigned criteria, such as gender, age, website, etc.

⁷ Note that *no* revenues are generated from processing reentrant fluid.

⁸ For $a, b \in \mathbb{R}^n$, $a \odot b = [a_1 \cdot b_1, \dots, a_n \cdot b_n]^T$.

Here we define an extremely simple control scheme for the problem we face. The scheme we propose will resolve a linear program similar to (2) at the times in \mathcal{T}_N with a certain “projected demand” based on conditions at the time of reoptimization.

3.1. The reoptimization scheme. For $(t, \Lambda, x) \in \mathbb{R}_+ \times \mathbb{R}_+^I \times \mathbb{R}^K$, define the linear program $\text{LP}(t, \Lambda, x)$ according to

$$\begin{aligned} & \max \sum_e p_e z_e \Lambda_{i(e)} \cdot (T - t) \\ & \text{subject to } \sum_e A_{k,e} z_e \Lambda_{i(e)} \cdot (T - t) \leq x_k^+ \quad \forall k, \\ & z \in \mathcal{Z}. \end{aligned}$$

Abusing notation, we will also denote the optimal value to this program by $\text{LP}(t, \Lambda, x)$. We will consider a reoptimization-based heuristic control policy $\{z_t^R\}$ defined so that

$$z_t^R = z_{d(t)}^R,$$

where $z_{iT/N}^R$ is any optimal solution to the linear program $\text{LP}(iT/N, \Lambda_{iT/N}, x_{iT/N})$. In words, this scheme assumes, at every point of time in $\hat{t} \in \mathcal{T}_N$ that the demand rate over the remaining time horizon will remain unchanged from $\Lambda_{\hat{t}}$, and employs the allocation rule that is optimal for such a scenario over the interval of time until the next re-solve. This procedure is summarized below

Reoptimization Heuristic

At each reoptimization interval $i = 0, \dots, N - 1$

1. measure demand rate $\Lambda_{iT/N}$;
2. obtain fractional allocation $z_{iT/N}^R \in \arg \max \text{LP}(iT/N, \Lambda_{iT/N}, x_{iT/N})$;
3. over the interval $[iT/N, (i + 1)T/N)$, allocate the demand $\int_{iT/N}^{(i+1)T/N} \Lambda_t dt$ according to $z_{iT/N}^R$.

Now define $J_{\{\Lambda_t\}}^R(x_0)$ as the revenues under the reoptimization heuristic under a specific sample path of demand, $\{\Lambda_t\}$, starting with inventory x_0 . In particular,

$$J_{\{\Lambda_t\}}^R(x_0) = \int_0^T \sum_e p_e \Lambda_{i(e),t} z_{e,t}^R \mathbf{1}_{\{I_{e,t}\}} dt,$$

where, as before, $I_{e,t} = \{x_{k,t} > 0 \forall k \text{ s.t. } A_{k,e} = 1\}$. We denote the total expected reward under the reoptimization policy assuming a starting inventory level x_0 by $J^R(x_0)$; $J^R(x_0) = E[J_{\{\Lambda_t\}}^R(x_0)]$.

3.2. An upper bound. Define $J_{\{\Lambda_t\}}^{\text{UB}}(x_0)$ as the optimal value of the (offline) optimization problem (2). That problem provides a useful upper bound on $J^{*,N}(x_0)$. In particular, we have

PROPOSITION 3.1. $E[J_{\{\Lambda_t\}}^{\text{UB}}(x_0)] \geq J^{*,N}(x_0)$ for all N .

The proof of this result can be found in the appendix—the result is natural; knowing realized demand a priori is in essence the best we can hope for.

3.3. Sample-path properties of the reoptimization policy and a lower bound. This section concerns two sample-path properties of the reoptimization policy. The first is simply a representation of $J_{\{\Lambda_t\}}^R(x_0)$ in terms of the optimal value of the linear programs solved along a sample path. The second simple but crucial property is a statement of “balanced” inventory consumption along a sample path.

Define $\Lambda_{d(t)}^{\min} \in \mathbb{R}_+^I$ so that $\Lambda_{i,d(t)}^{\min} = \min\{\Lambda_{i,t} : d(t) \leq t < d(t) + N/T\}$. Similarly define $\Lambda_{d(t)}^{\max}$. We then have the following lemma:

LEMMA 3.1.

$$J_{\{\Lambda_t\}}^R(x_0) \geq \frac{T}{N} \sum_{j=0}^{N-1} \left(\frac{\text{LP}(jT/N, \Lambda_{jT/N}, x_{jT/N})}{T - jT/N} - C \left(\sum_e p_e \right) \sum_e (\Lambda_{i(e),jT/N}^{\max} - \Lambda_{i(e),jT/N}^{\min}) \right),$$

where C is some constant dependent only on the quantities $A_{k,e}$, $k \in 1, \dots, K$, $e \in E$.

The proof of this lemma is relatively routine and can be found in the appendix. The continuity of the sample paths of $\{\Lambda_t\}$ yield the following result as an easy corollary, whose proof is also in the appendix.

COROLLARY 3.1.

$$\liminf_N J_{\{\Lambda_t\}}^R(x_0) \geq \liminf_N \frac{T}{N} \sum_{j=0}^{N-1} \frac{\text{LP}(jT/N, \Lambda_{jT/N}, x_{jT/N})}{T - jT/N}.$$

We next demonstrate a crucial sample path property that will eventually allow us to relate the linear programs solved at each opportunity to reoptimize to the initial (offline) LP solved at $t = 0$. The property is natural and simple to derive, but quite powerful in its application.

LEMMA 3.2 (BALANCING). *For every sample path of Λ_t ,*

$$x_{k, nT/N} \geq x_{k, 0} \frac{N - n}{N} + \sum_{j=0}^{n-1} \sum_e \frac{T}{N} A_{k, e} (\Lambda_{i(e), jT/N} - \Lambda_{i(e), jT/N}^{\max})$$

for all $n < N$.

PROOF. We proceed by induction. The claim is trivially true for $n = 0$. Assume the claim true for $n = l - 1 < N - 1$ and observe that

$$\begin{aligned} x_{k, (l+1)T/N} &\geq x_{k, lT/N} - \sum_e A_{k, e} \left(\int_{lT/N}^{(l+1)T/N} \Lambda_{i(e), t} dt \right) z_{e, lT/N}^R \\ &\geq x_{k, lT/N} - \sum_e A_{k, e} \left(\int_{lT/N}^{(l+1)T/N} \Lambda_{i(e), lT/N}^{\max} dt \right) z_{e, lT/N}^R \\ &= x_{k, lT/N} - \sum_e \frac{T}{N} A_{k, e} \Lambda_{i(e), lT/N} z_{e, lT/N}^R + \sum_e \frac{T}{N} A_{k, e} (\Lambda_{i(e), lT/N} - \Lambda_{i(e), lT/N}^{\max}) z_{e, lT/N}^R \\ &\geq x_{k, lT/N} \left(1 - \frac{1}{N - l} \right) + \sum_e \frac{T}{N} A_{k, e} (\Lambda_{i(e), lT/N} - \Lambda_{i(e), lT/N}^{\max}) z_{e, lT/N}^R. \end{aligned} \tag{3}$$

The final inequality follows from the fact that $z_{lT/N}^R$ is a feasible solution to $\text{LP}(lT/N, \Lambda_{lT/N}, x_{lT/N})$. But the induction hypothesis yields

$$x_{k, lT/N} \geq x_{k, 0} \frac{N - l}{N} + \sum_{j=0}^{l-1} \sum_e \frac{T}{N} A_{k, e} (\Lambda_{i(e), jT/N} - \Lambda_{i(e), jT/N}^{\max}),$$

and substituting in the final inequality of (3) then yields

$$x_{k, (l+1)T/N} \geq x_{k, 0} \frac{N - (l + 1)}{N} + \sum_{j=0}^l \sum_e \frac{T}{N} A_{k, e} (\Lambda_{i(e), jT/N} - \Lambda_{i(e), jT/N}^{\max}).$$

Induction on l completes the proof. \square

Lemma 3.2 yields as a corollary the following result whose proof may be found in the appendix.

COROLLARY 3.2.

$$\liminf_N \frac{T}{N} \sum_{j=0}^{N-1} \frac{\text{LP}(jT/N, \Lambda_{jT/N}, x_{jT/N})}{T - jT/N} \geq \liminf_N \frac{T}{N} \sum_{j=0}^{N-1} \frac{\text{LP}(jT/N, \Lambda_{jT/N}, x_0(N - j)/N)}{T - jT/N}.$$

Corollaries 3.1 and 3.2 together imply the following lower bound on revenues under the reoptimization heuristic:

THEOREM 3.1.

$$\liminf_N J_{\{\Lambda_t\}}^R(x_0) \geq \liminf_N \frac{T}{N} \sum_{j=0}^{N-1} \frac{\text{LP}(jT/N, \Lambda_{jT/N}, x_0(N - j)/N)}{T - jT/N}.$$

3.4. Properties of the re-solved LP and a decomposition. Here we briefly develop an “expansion” of the optimal value $LP(t, \Lambda, x)$ that is separable in the components of Λ . This expansion will serve us in our performance analysis and, in particular, will be crucial in analyzing an inherently multidimensional system via a single-dimensional analysis. We begin with some definitions.

For every $i \in \{1, \dots, I\}$ and tuples $(t, \Lambda, x) \in \mathbb{R}_+ \times \mathbb{R}_+^I \times \mathbb{R}_+^K$, let $z(t, \Lambda, x) \in \mathbb{R}^E$ denote an optimal solution to $LP(t, \Lambda, x)$, and define the functions $f_i^{t, \Lambda, x}: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ according to

$$f_i^{t, \Lambda, x}(w) = \sum_{e: i(e)=i} p_e \min\{\Lambda_i, w\} z_e(t, \Lambda, x).$$

We catalog a few properties of these functions that will serve us well in the sequel; these properties are proved in the appendix.

LEMMA 3.3. *For every $i \in \mathcal{I}$, and $(t, \Lambda, x) \in \mathbb{R}_+ \times \mathbb{R}_+^I \times \mathbb{R}_+^K$, the function $f(w) \triangleq f_i^{t, \Lambda, x}(w)$ satisfies the following properties:*

- (i) $f(0) = 0$ and f is continuous and nondecreasing;
- (ii) $f(\cdot)$ is concave;
- (iii) for $w, v > 0$, $f(w)/f(v) \geq \min\{w/v, 1\}$;
- (iv) $(1/x) \int_0^x f(w) dw \leq f(x/2)$.

The utility of the functions $f_i^{t, \Lambda, x}$ lies in the fact that they allow us to construct useful approximations to $LP(t, \Lambda, x)$ that are separable in the components of Λ . This is made precise by the following result whose proof is deferred to the Appendix:

LEMMA 3.4. *For any $(t, u, x) \in \mathbb{R}_+ \times \mathbb{R}_+^I \times \mathbb{R}_+^K$, and an arbitrary $\Lambda \in \mathbb{R}_+^I$, we have*

$$LP(t, u, x) \geq (T - t) \sum_i f_i^{(t, \Lambda, x)}(u_i)$$

with equality for $u = \Lambda$.

3.5. Performance guarantees for the reoptimization scheme. With the lower and upper bounds developed thus far, we are finally in a position to present performance guarantees for our approach for processes where $\lambda_{i,t} = \lambda_i$.

THEOREM 3.2. *For demand processes $\{\Lambda_t\}$ satisfying Assumption 2.1, and with $\lambda_{i,t} = \lambda_i$ for all i, t , we have, assuming an initial inventory of x_0*

$$\liminf_N \frac{\mathbb{E}[J_{\{\Lambda_t\}}^R(x_0)]}{\mathbb{E}[J_{\{\Lambda_t\}}^{UB}(x_0)]} \geq 0.342.$$

PROOF. We know by Fatou’s Lemma that

$$\liminf_N \frac{\mathbb{E}[J_{\{\Lambda_t\}}^R(x_0)]}{\mathbb{E}[J_{\{\Lambda_t\}}^{UB}(x_0)]} \geq \frac{\mathbb{E}[\liminf_N J_{\{\Lambda_t\}}^R(x_0)]}{\mathbb{E}[J_{\{\Lambda_t\}}^{UB}(x_0)]}.$$

We will proceed in turn to bound the numerator and denominator. We then have that

$$\begin{aligned} \liminf_N J_{\{\Lambda_t\}}^R(x_0) &\geq \liminf_N \frac{T}{N} \sum_{j=0}^{N-1} \frac{LP(0, \Lambda_{jT/N}, x_0)}{T} \\ &\geq \liminf_N \frac{T}{N} \sum_{j=0}^{N-1} \sum_i f_i^{(0, \hat{\Lambda}, x_0)}(\Lambda_{i, jT/N}) \\ &= \sum_i \int_0^T f_i^{(0, \hat{\Lambda}, x_0)}(\Lambda_{i,t}) dt, \end{aligned} \tag{4}$$

where $\hat{\Lambda} \in \mathbb{R}_+^I$ is arbitrary. The first inequality is a consequence of Corollaries 3.1 and 3.2 and the fact that $LP(t, \Lambda, x(1 - t/T)) = ((T - t)/T)LP(0, \Lambda, x)$ for $0 \leq t \leq T$, and arbitrary $\Lambda \in \mathbb{R}_+^I$, $x \in \mathbb{R}_+^K$. The second inequality is Lemma 3.4. The final equality follows from the continuity of $f_i^{(0, \hat{\Lambda}, x_0)}(\cdot)$ (Lemma 3.3, property 1) and the continuity of $\Lambda_{i,t}$ in t .

Next, we have that

$$\mathbb{E}[J_{\{\Lambda_t\}}^{\text{UB}}(x_0)] = \mathbb{E}[\text{LP}(0, \bar{\Lambda}, x_0)] \leq \text{LP}(0, \mathbb{E}[\bar{\Lambda}], x_0) = \sum_i T f_i^{(0, \mathbb{E}[\bar{\Lambda}], x_0)}(\mathbb{E}[\bar{\Lambda}_i]), \quad (5)$$

where we define $\bar{\Lambda} \in \mathbb{R}_+^I$ according to $\bar{\Lambda}_i = (1/T) \int_0^T \Lambda_{i,t} dt$. The inequality above is Jensen’s inequality (since $\text{LP}(t, u, x)$ is concave in u). The final equality is Lemma 3.4. Now, from (4) and (5), it follows that

$$\liminf_N \frac{\mathbb{E}[J_{\{\Lambda_t\}}^{\text{R}}(x_0)]}{\mathbb{E}[J_{\{\Lambda_t\}}^{\text{UB}}(x_0)]} \geq \frac{\sum_i \mathbb{E}[\int_0^T (1/T) f_i^{(0, \mathbb{E}[\bar{\Lambda}], x_0)}(\Lambda_{i,t}) dt]}{\sum_i f_i^{(0, \mathbb{E}[\bar{\Lambda}], x_0)}(\mathbb{E}[\bar{\Lambda}_i])} \geq \min_i \frac{\mathbb{E}[\int_0^T (1/T) f_i^{(0, \mathbb{E}[\bar{\Lambda}], x_0)}(\Lambda_{i,t}) dt]}{f_i^{(0, \mathbb{E}[\bar{\Lambda}], x_0)}(\mathbb{E}[\bar{\Lambda}_i])} \geq 0.342,$$

where the final inequality is the estimate derived in Lemma 2.1. \square

This guarantee is remarkable in that it is *uniform* over a broad class of demand processes. In particular, the guarantee places no restrictions whatsoever on the volatility of these processes nor on their correlation or autocorrelation structures. Of course, in the event that the volatility of the underlying process were small, one expects even better performance from this very natural algorithm, and indeed a proof essentially identical to this one yields the following theorem:

THEOREM 3.3. *Consider demand processes $\{\Lambda_t\}$ satisfying Assumption 2.1, with $\lambda_t = \lambda$ for all t . In addition assume that $\sigma_t/\lambda \leq \sqrt{2\pi}B$. We then have, assuming an initial inventory of x_0*

$$\liminf_N \frac{\mathbb{E}[J_{\{\Lambda_t\}}^{\text{R}}(x_0)]}{\mathbb{E}[J_{\{\Lambda_t\}}^{\text{UB}}(x_0)]} \geq \max \left\{ 0.342, \frac{1}{1+B} - \frac{B}{1+B} (\exp(-1/4\pi B^2) + 0.853) \right\}.$$

The proof of this above theorem proceeds identically to that of Theorem 3.2 with the exception that, in the very final inequality of that proof, we use the general bound given in Remark 2.1.

Theorems 3.2 and 3.3 together establish a strong statement about the robustness of our reoptimization heuristic. In particular, these theorems establish that with frequent reoptimization, this natural heuristic attains the ability to compete with a clairvoyant with perfect knowledge of the sample paths of the demand process, irrespective of the volatility of that process. Simultaneously, in the event that the underlying process was not volatile at all the same scheme is essentially optimal.

3.5.1. Essential properties of the rate process and potential generalizations. The essence of Theorem 3.2 is that it reduces the performance analysis of our reoptimization scheme to the analysis of *simple* properties of *marginals* of the rate process. As it happens, these properties can be tractably quantified for moving average processes (which as we established earlier are a particularly interesting family of processes). However, it is possible that such guarantees might be established for other classes of processes, and to this end we explicitly isolate the property of the rate process that drives the bound of Theorem 3.2 and then present an example of a process that is not a moving average process but admits a constant factor guarantee.

Notice that the proof of Theorem 3.2 yields

$$\liminf_N \frac{\mathbb{E}[J_{\{\Lambda_t\}}^{\text{R}}(x_0)]}{\mathbb{E}[J_{\{\Lambda_t\}}^{\text{UB}}(x_0)]} \geq \min_i \frac{\mathbb{E}[(1/T) \int_0^T f_i(\Lambda_{i,t}) dt]}{f_i(\mathbb{E}[(1/T) \int_0^T \Lambda_{i,t} dt])}$$

for some set of functions f_i , each satisfying Lemma 3.3, and *any* rate process with continuous sample paths. We were able to come up with a uniform lower bound to this quantity for arbitrary, nonstationary moving average processes.

The analysis of this quantity is substantially simplified in the case of stationary processes. In particular, consider processes that are *stationary* and whose marginals have finite expectation. If this were the case, we have

$$\begin{aligned} \frac{\mathbb{E}[(1/T) \int_0^T f_i(\Lambda_{i,t}) dt]}{f_i(\mathbb{E}[(1/T) \int_0^T \Lambda_{i,t} dt])} &\geq \mathbb{E} \left[\frac{1}{T} \int_0^T \min \left\{ \frac{\Lambda_{i,t}}{\mathbb{E}[(1/T) \int_0^T \Lambda_{i,t} dt]}, 1 \right\} dt \right] \\ &= \mathbb{E} \left[\frac{1}{T} \int_0^T \min \left\{ \frac{\Lambda_{i,t}}{\mathbb{E}[\Lambda_{i,t}]}, 1 \right\} dt \right] \\ &= \mathbb{E} \left[\min \left\{ \frac{\Lambda_{i,0}}{\mathbb{E}[\Lambda_{i,0}]}, 1 \right\} \right], \end{aligned}$$

where we have used Lemma 3.3 for the first inequality, and the stationarity of the process with Fubini’s theorem for the next two equalities. Notice that this bound has a natural interpretation. It measures, in a sense, the asymmetry of the marginals of the process under consideration. In general, this quantity can be arbitrarily small (consider, for instance, a suitable two-point distribution). However, as witnessed by the moving average family of processes, there are potentially *large* families of stochastic processes for which this quantity is uniformly bounded for all processes within the family.

Here we give another example. Imagine that each of the I marginals are described by a *Cox-Ingersoll-Ross* (CIR) process. A CIR process is a nonnegative, mean-reverting process and is perhaps the best known example of an *affine* process. It is typically used to model the behavior of nonnegative quantities such as interest rates or (more recently) arrival rates to a queueing system (see Besbes and Maglaras [8]). The process is defined as the solution of the stochastic differential equation

$$d\Lambda_{i,t} = \theta_i(\lambda_i - \Lambda_{i,t}) dt + \sigma_i \sqrt{\Lambda_{i,t}} dZ_{it} \quad \text{for } \theta_i, \lambda_i, \sigma_i > 0.$$

When $2\theta_i\lambda_i > \sigma_i^2$ (the regime typically considered in any modeling with this process), the CIR process is strictly positive and has a stationary distribution. This stationary distribution is Gamma (as opposed to Gaussian for moving average processes) with shape parameter $a \triangleq 2\theta_i\lambda_i/\sigma_i^2$ and scale parameter $\sigma_i^2/2\theta_i$. Consequently, for *stationary* CIR processes we may compute (see Farias and Van Roy [14]),

$$\mathbb{E} \left[\min \left\{ \frac{\Lambda_{i,0}}{\mathbb{E}[\Lambda_{i,0}]}, 1 \right\} \right] = 1 - \frac{\Gamma(a+1, a)}{\Gamma(a+1)} + \frac{\Gamma(a, a)}{\Gamma(a)} \geq 1 - \frac{1}{e}$$

for $a \geq 1$. This yields the following theorem:

THEOREM 3.4. *For demand processes whose marginals are stationary (but otherwise arbitrary) CIR processes, we have assuming an initial inventory of x_0*

$$\liminf_N \frac{\mathbb{E}[J_{\{\Lambda_i\}}^R(x_0)]}{\mathbb{E}[J_{\{\Lambda_i\}}^{UB}(x_0)]} \geq 1 - \frac{1}{e}.$$

Of course, the class of processes here, while allowing for arbitrary volatility, is nowhere as large or perhaps interesting as the moving average processes we have focused on, but the result provides a sense of the generalizability of the analysis.

3.6. The impact of reoptimization frequency. Our performance guarantees thus far call for frequent reoptimization (i.e., large N). At this juncture we ask two questions:

(i) Is a large value of N —and the implicit demand forecast updating—truly necessary for good performance of our scheme?

(ii) If large values of N are indeed necessary, what impact does a finite number of reoptimizations have on the performance guarantee of Theorem 3.2?

We will answer the first question in the affirmative by providing a sequence of problems where infrequent reoptimization results in poor (in fact, arbitrarily poor) performance. As for the second question, we will use a well-known result on the modulus of continuity for the sample paths of Brownian motion to characterize the error because of finite reoptimization.

Throughout this section, we make the dependence of the revenue on the number of reoptimization intervals explicit through the notation $J_{\{\Lambda_i\}}^{R, N}$ (in place of simply $J_{\{\Lambda_i\}}^R$). To answer the first question and demonstrate the importance of resolving, we will describe a sequence of problems, indexed by T , and show that if we choose to not reoptimize (and thereby not adjust forecasts) under our scheme, such a choice can grow arbitrarily sub-optimal as T grows large. This will show that reoptimization plays a dramatic role in the actual performance of our scheme. We next describe this sequence of problems.

EXAMPLE 3.1. Consider an allocation problem with two sources and a single sink. We have one resource type so that $K = 1$, and set $A_{1,(1,1)} = A_{1,(2,1)} = 1$, while we set $x_0 = T$. Our time horizon is T , and we set $p_{(1,1)} = 1$ and $p_{(2,1)} = T^{1/3+\epsilon}$ where $\epsilon > 0$ is some constant. The demand process is described as follows: $\Lambda_{1,t} = (\frac{1}{2} + \sqrt{2\pi}W_t)^+$, where W_t is standard Brownian motion and $\Lambda_{2,t} = \frac{1}{2}$. Notice that this demand process satisfies Assumption 2.1.

We then have the following result, whose proof may be found in the appendix.

PROPOSITION 3.2. *For the problem described in Example 3.1, we have*

$$\frac{J^{R,1}(x_0)}{J^{UB}(x_0)} = O(T^{-1/3}).$$

Now, in contrast, Theorem 3.2 implies that

$$\liminf_N \frac{J^{R,N}(x_0)}{J^{UB}(x_0)} = \Theta(1).$$

Contrasting these two results demonstrates the importance of frequent reoptimization in the context of using our allocation scheme.

In light of the result, we can move on to answering the next logical question, which asks us to establish the effect of a finite number of reoptimizations, or more specifically a “rate” for the limit infimum in Theorem 3.2. We will establish such a rate for moving average demand rate processes. Doing this will require the following principal technical tool, which is a *global* modulus of continuity for sample paths of moving average processes (see, for instance Karatzas and Shreve [17]).

THEOREM 3.5 (LEVY’S MODULUS OF CONTINUITY). *Let $\bar{\Lambda}_t$ be a moving average process. Then, almost surely,*

$$\limsup_{h \rightarrow 0} \sup_{0 \leq t \leq T-h} \frac{|\bar{\Lambda}_{t+h} - \bar{\Lambda}_t|}{\sqrt{2h \log 1/h}} \leq \phi(0).$$

Note that Levy’s theorem is typically stated for standard Brownian motion where the limit above can be shown to exist, and is equal to one. This result is, in fact, a simple corollary to that theorem. Roughly, the theorem can be interpreted as stating that $\sup_{0 \leq t \leq T-h} |\bar{\Lambda}_t - \bar{\Lambda}_{t+h}| = O(\sqrt{h \log 1/h})$. We will now employ this theorem to prove a performance guarantee for the reoptimization scheme allowing only a finite number of reoptimizations.

THEOREM 3.6. *Assume the demand rate process $\{\Lambda_t\}$ satisfies $\Lambda_{i,t} = (\bar{\Lambda}_{i,t})^+$, where $\bar{\Lambda}_t$ is a multivariate moving average process. Define $\sigma_i \triangleq \phi_i(0)$. Then,*

$$E[J_{\{\Lambda_t\}}^{R,N}(x_0)] \geq 0.342 E[J_{\{\Lambda_t\}}^{UB}(x_0)] - \Delta(N),$$

where $\Delta(N)$ satisfies

$$\limsup_N \frac{\Delta(N)}{\sqrt{\log N/N}} \leq C' \sum_i \sigma_i.$$

The constant C' can be specified independent of the demand process.

The constant C' is derived explicitly in the proof of the theorem (which can be found in the appendix). It depends solely on the quantities p, A, K, T , and I . As such, the result can loosely be interpreted as stating that a finite number of reoptimizations introduces an *additive* error that behaves roughly like $O(\sum_i \sigma_i \sqrt{\log N/N})$. This has an interesting interpretation: the additive error component grows at most linearly with the volatility of the process. However, with sufficiently frequent reoptimization (as specified by the rate in the theorem), this additive error can be made arbitrarily small. Put another way, frequent reoptimization is particularly valuable when the demand rate process in question is highly volatile.

4. A heuristic that uses forecasts: β reoptimization. In this section, we develop and analyze a heuristic that has access to a deterministic forecast of demand evolution. In the context of our assumed model of stochastic demand processes, we will assume knowledge of λ_t . The heuristic we develop is closely related to the reoptimization heuristic, and will be competitive for processes where λ_t is not necessarily constant.

Imagine a scenario wherein the demand process had no noise—in particular, $\sigma_t = 0$. In this event, any optimal solution to $LP(0, \int_0^T \lambda_t dt, x_0)$ constitutes an optimal (and static) allocation rule. Of course, if $\sigma_t > 0$, this is not the case, and so the policy we propose will entail a careful convex combination of this “deterministic” policy with a reoptimization policy analogous to that studied in the previous section.

Informally speaking, our heuristic will approximately simulate the following allocation over time:

(i) Split the total demand $\{\Lambda_t\}$ and resource capacity x_0 into two systems. System 1 (the reoptimization system) sees the demand process $\{(1 - \beta)\Lambda_t\}$ and begins with inventory $(1 - \beta)x_0$. System 2 (the deterministic system) sees the demand process $\{\beta\Lambda_t\}$ and begins with inventory βx_0 .

(ii) Apply the reoptimization policy (of the last section) $\{z_t^R\}$ to the reoptimization system.

(iii) Apply the deterministic policy (to be defined momentarily) $\{z_t^D\}$ to the deterministic system.

This section will establish uniform performance guarantees for the policy (loosely); these guarantees will implicitly identify an oblivious choice of β that is “good.” In addition, we will provide a sketch of potential improvements to the policy and a guideline on how to tune β given further information about the demand process. In broad steps, the analysis will proceed along the lines of the following roadmap:

(i) Establish a performance guarantee for the reoptimization policy with respect to an upper bound on revenues for the reoptimization system. This will utilize our previous analysis. See Lemma 4.2.

(ii) Establish a performance guarantee for the deterministic policy with respect to an upper bound on revenues for the deterministic system. See Lemma 4.3.

(iii) Establish a relationship between optimal revenues for the reoptimization and deterministic systems with optimal revenues for the original problem. See Lemma 4.4.

(iv) Using the above three steps, compute a performance guarantee for the overall scheme. This guarantee will be a function of β , which we may then optimize. See Theorem 4.1.

We next define our policy formally.

4.1. The β reoptimization scheme. We first construct a few auxiliary processes. In particular, define the reoptimization inventory process \hat{x}_t^R according to

$$\hat{x}_{k,t}^R = \hat{x}_{k,0}^R - \int_0^t \mathbf{1}_{\hat{x}_{k,t}^R > 0} \sum_e \Lambda_{i(e),t} A_{k,e} \hat{z}_{e,d(t)}^R dt,$$

where $\hat{x}_0^R = x_0$ and $\hat{z}_{iT/N}^R$ is any optimal solution to $\text{LP}(iT/N, \Lambda_{iT/N}, \hat{x}_{iT/N}^R)$. In words, \hat{x}_t^R is the inventory process obtained if one employs the allocation policy $z_t = \hat{z}_{d(t)}^R$.

Next, let us denote by z^D , an optimal solution to $\text{LP}(0, \int_0^T \lambda_t dt, x_0)$. Define the deterministic policy $\{\hat{z}_t^D\}$ according to $\hat{z}_t^D \triangleq (\lambda_{d(t)}/\Lambda_{d(t)} \wedge 1)z^D$. Note that $\{\hat{z}_t^D\}$ is an optimal allocation policy in the event $\sigma_t = 0$.

The control we propose, denoted $\{z_t^{R-\beta}\}$, is defined according to

$$z_t^{R-\beta} = (1 - \beta)\hat{z}_{d(t)}^R + \beta\hat{z}_{d(t)}^D.$$

β Reoptimization heuristic

1. Compute z^D .
2. At each reoptimization interval $i = 0, \dots, N - 1$
 - a. measure demand rate $\Lambda_{iT/N}$;
 - b. obtain allocation $\hat{z}_{iT/N}^R \in \arg \max \text{LP}(iT/N, \Lambda_{iT/N}, \hat{x}_{iT/N}^R)$;
 - c. over the interval $[iT/N, (i + 1)T/N)$, allocate the demand according to $(1 - \beta)\hat{z}_{iT/N}^R + \beta\hat{z}_{iT/N}^D$.

4.2. Preliminary sample path properties. Here we identify several sample path properties for the β reoptimized policy. Define the reoptimized revenue under the β reoptimized policy according to

$$J_{S, \{\Lambda_t\}}^R = \int_0^T \sum_e p_e \hat{z}_{e,d(t)}^R \Lambda_{i(e),t} \mathbb{1}_{\{\hat{x}_{e,t}^R > 0\}} dt,$$

where $\hat{x}_{e,t}^R = \{\hat{x}_{k,t}^R > 0 \ \forall k \text{ s.t. } A_{k,e} > 0\}$. Similarly define the deterministic revenue under the β reoptimized policy according to

$$J_{D, \{\Lambda_t\}}^R = \int_0^T \sum_e p_e \hat{z}_{e,d(t)}^D \Lambda_{i(e),t} dt.$$

Given the definitions of \hat{x}_t^R and $\hat{z}_{jT/N}^R$, Theorem 3.1 immediately implies

$$\liminf_N J_{S, \{\Lambda_t\}}^R(x_0) \geq \liminf_N \frac{T}{N} \sum_{j=0}^{N-1} \frac{\text{LP}(jT/N, \Lambda_{jT/N}, \hat{x}_0^R(N-j)/N)}{T - jT/N}. \tag{6}$$

Moreover, we have

$$\begin{aligned} J_{D, \{\Lambda_t\}}^R(x_0) &= \sum_e p_e \int_0^T \Lambda_{i(e),t} \hat{z}_{e,t}^D dt \geq \frac{T}{N} \sum_{j=0}^{N-1} \sum_e p_e \Lambda_{i(e),jT/N}^{\min} \hat{z}_{e,jT/N}^D \\ &= \frac{T}{N} \sum_{j=0}^{N-1} \left(\sum_e p_e \Lambda_{i(e),jT/N} \hat{z}_{e,jT/N}^D - \sum_e p_e (\Lambda_{i(e),jT/N} - \Lambda_{i(e),jT/N}^{\min}) \right) \end{aligned}$$

so that

$$\liminf_N J_{D, \{\Lambda_t\}}^R(x_0) \geq \liminf_N \frac{T}{N} \sum_{j=0}^{N-1} \sum_e p_e \Lambda_{i(e), jT/N} \hat{z}_{e, jT/N}^D. \quad (7)$$

Finally, we have the following result that decomposes the revenues under the β reoptimized policy into the reoptimized and deterministic revenue components; the proof may be found in the appendix.

LEMMA 4.1.

$$\liminf_N E[J_{\{\Lambda_t\}}^{R-\beta}(x_0)] \geq \liminf_N (1 - \beta) E[J_{S, \{\Lambda_t\}}^R(x_0)] + \liminf_N \beta E[J_{D, \{\Lambda_t\}}^R(x_0)].$$

REMARK 4.1. We note that although we will not need this fact, the inequality established in the result in Lemma 4.1 also holds on a sample path basis; In particular, we can show

$$\liminf_N J_{\{\Lambda_t\}}^{R-\beta}(x_0) \geq \liminf_N (1 - \beta) J_{S, \{\Lambda_t\}}^R(x_0) + \liminf_N \beta J_{D, \{\Lambda_t\}}^R(x_0).$$

4.3. Performance analysis. This section establishes uniform performance guarantees on the performance of the β reoptimization scheme. Using the decomposition arrived at it the previous section, this guarantee is arrived at by deriving appropriate uniform guarantees on the reoptimized and deterministic revenue terms. The former guarantee is essentially obtained via Theorem 3.2 whereas the latter requires a new argument. Our arguments will require one extra technical assumption on the demand process $\{\Lambda_t\}$:

ASSUMPTION 4.1. For all i , $E[\max_{t \in [0, T]} \Lambda_{i,t}] < \infty$.

We begin with an analysis of the reoptimized revenue component, $J_{S, \{\Lambda_t\}}^R$. The proof is essentially a corollary to Theorem 3.2 and deferred to the appendix.

LEMMA 4.2.

$$\liminf_N \frac{E[J_{S, \{\Lambda_t\}}^R(x_0)]}{E[J_{\{\Lambda_t, -\lambda_t\}^+}^{UB}(x_0)]} \geq 0.342.$$

Our next result provides an analysis of the deterministic component of revenues under the β reoptimization scheme, $J_{D, \{\Lambda_t\}}^R$.

LEMMA 4.3.

$$\liminf_N \frac{E[J_{D, \{\Lambda_t\}}^R(x_0)]}{J_{\{\lambda_t\}}^{UB}(x_0)} \geq \frac{1}{2}.$$

PROOF. Now, we have

$$\begin{aligned} \liminf_N E[J_{D, \{\Lambda_t\}}^R(x_0)] &\geq E[\liminf_N J_{D, \{\Lambda_t\}}^R(x_0)] \\ &\geq E\left[\liminf_N \frac{T}{N} \sum_{j=0}^{N-1} \sum_e p_e \Lambda_{i(e), jT/N} \hat{z}_{e, jT/N}^D\right] \\ &= \lim_N E\left[\frac{T}{N} \sum_{j=0}^{N-1} \sum_e p_e \Lambda_{i(e), jT/N} \hat{z}_{e, jT/N}^D\right] \\ &\geq \liminf_N E\left[\frac{T}{N} \sum_{j=0}^{N-1} \sum_e p_e \mathbb{1}_{\{\Lambda_{i(e), jT/N} \geq \lambda_{i(e), jT/N}\}} \lambda_{i(e), jT/N} \hat{z}_{e, jT/N}^D\right] \\ &= \lim_N \frac{T}{N} \sum_{j=0}^{N-1} \sum_e p_e \frac{1}{2} \lambda_{i(e), jT/N} \hat{z}_{e, jT/N}^D \\ &= \frac{1}{2} J_{\{\lambda_t\}}^{UB}(x_0). \end{aligned}$$

The first inequality is Fatou's Lemma and the second follows from (7). The second equality follows from the dominated convergence theorem: in particular, observe that

$$\lim_N \frac{T}{N} \sum_{j=0}^{N-1} \sum_e p_e \Lambda_{i(e), jT/N} \hat{z}_{e, jT/N}^D$$

exists by the definition of $\hat{z}_{e, jT/N}^D$ and the continuity of $\Lambda_{i,t}$ and $\lambda_{i,t}$. Further,

$$\sum_{j=0}^{N-1} \sum_e p_e \Lambda_{i(e), jT/N} \hat{z}_{e, jT/N}^D \leq E \sum_i \max_{t \in [0, T]} \Lambda_{i,t},$$

which was assumed finite by Assumption 4.1. \square

Before moving on to our approximation guarantee, we establish one last fact (the proof is in the appendix):

LEMMA 4.4.

$$J_{\{\lambda_t\}}^{UB}(x_0) + E[J_{\{(\bar{\lambda}_t - \lambda_t)^+\}}^{UB}(x_0)] \geq E[J_{\{\Lambda_t\}}^{UB}(x_0)].$$

We are now in a position to provide a uniform performance guarantee for the β reoptimized scheme. In particular, we have

THEOREM 4.1.

$$\liminf_N \frac{E[J_{\{\lambda_t\}}^{R-\beta}(x_0)]}{E[J_{\{\lambda_t\}}^{UB}(x_0)]} \geq (1 - \beta)0.342 \wedge \beta 0.5.$$

PROOF. We have

$$\begin{aligned} \liminf_N \frac{E[J_{\{\lambda_t\}}^{R-\beta}(x_0)]}{E[J_{\{\lambda_t\}}^{UB}(x_0)]} &\geq \frac{\liminf_N (1 - \beta) E[J_{S, \{\lambda_t\}}^R(x_0)] + \liminf_N \beta E[J_{D, \{\lambda_t\}}^R(x_0)]}{E[J_{\{(\bar{\lambda}_t - \lambda_t)^+\}}^{UB}(x_0)] + J_{\{\lambda_t\}}^{UB}(x_0)} \\ &\geq \frac{(1 - \beta)0.342 E[J_{\{(\bar{\lambda}_t - \lambda_t)^+\}}^{UB}(x_0)] + \beta 0.5 J_{\{\lambda_t\}}^{UB}(x_0)}{E[J_{\{(\bar{\lambda}_t - \lambda_t)^+\}}^{UB}(x_0)] + J_{\{\lambda_t\}}^{UB}(x_0)} \\ &= (1 - \beta)0.342 \wedge \beta 0.5. \end{aligned}$$

The first inequality follows from Lemmas 4.4 and 4.1. The second inequality follows from Lemmas 4.3 and 4.2. \square

Before closing this section, we remark on the implications of the result and several issues related to implementing the β reoptimized scheme in practice:

(i) Choosing β . Optimizing the bound in Theorem 4.1 suggests setting $\beta \sim 0.4$. This is an oblivious choice of β that results in a uniform performance guarantee.

(ii) With further knowledge about the demand process one might be able to do better. For instance, an estimate of the relative values of $J_{\{\lambda_t\}}^{UB}(x_0)$ and $E[J_{\{(\bar{\lambda}_t - \lambda_t)^+\}}^{UB}(x_0)]$ (which can both be computed easily if the model for Λ_t were known) suggests a better selection rule: set $\beta = 1$ if $J_{\{\lambda_t\}}^{UB}(x_0) > 0.684 E[J_{\{(\bar{\lambda}_t - \lambda_t)^+\}}^{UB}(x_0)]$ and set $\beta = 0$ otherwise. A practical guideline is to set $\beta = 1$ when we believe there is little to no volatility in the demand process, and set $\beta = 0$ otherwise. Our numerical experiments (described in the following section) seem to suggest that, for natural problem instances, volatility quickly drowns out the value of forecasts, and hence $\beta = 1$ tends to perform best in most instances where demand uncertainty is present.

(iii) A practical improvement to the algorithm that does not alter the guarantees moves inventory made available because of underutilization by the deterministic system to the stochastic system. This compensates for the lack of inventory sharing between the reoptimized and deterministic systems. In particular, while leaving all of the details of the algorithm unchanged, we define the dynamic for \hat{x}_t^R according to

$$\hat{x}_{k,t}^R = \hat{x}_{k,0}^R - \int_0^t \mathbf{1}_{\hat{x}_{k,t}^R > 0} \sum_e \Lambda_{i(e),t} A_{k,e} \hat{z}_{e,d(t)}^R dt + \int_0^t \sum_e (\lambda_{i(e),t} z_e^D - \Lambda_{i(e),t} \hat{z}_{e,t}^D) dt.$$

5. Experiments. We focus our experiments on instances of the ad display problem. We present results for two sets of experiments. The first set consists of synthetic instances; the purpose of this set of experiments is to gauge performance across a variety of parameter regimes. The second set of experiments is derived from an actual allocation problem and (real) traffic from an ad network.

5.1. A generative family of instances. We characterize problem instances along two dimensions, namely,

(i) *Load factor.* We define load factor as the quantity

$$LF \triangleq \frac{E[\sum_i \int_0^T \bar{\Lambda}_{i,t} dt]}{\sum_k x_{k,0}}.$$

This is a natural measure of the scarcity (or abundance) of a resource relative to demand.

(ii) *Coefficient of variation (CV)*. We measure the relative volatility in demand via the quantity

$$CV \triangleq \frac{\sqrt{\text{Var}[\sum_i \int_0^T \bar{\Lambda}_{i,t} dt]}}{\mathbb{E}[\sum_i \int_0^T \bar{\Lambda}_{i,t} dt]}.$$

We consider the following generative family of instances:

(i) *Topology*. We set $I = 30$ and $A = 30$. An edge connecting a given source i to a given sink a exists independently with probability 0.1. In essence, this prevents a scenario where a given unit of traffic can be used by essentially all sinks.

(ii) *Resources, prices, and horizon*. We set $x_{k,0} = 100$. The price associated with a given edge is generated according to an independent uniform distribution on $[0, 100]$. We set $T = 1$.

(iii) *Demand*. We use an Ornstein-Uhlenbeck process to generate Λ_t . In particular, we set

$$\bar{\Lambda}_t = \lambda_t + \sigma \int_0^t e^{(s-t)} dZ_s,$$

where Z_s is standard I dimensional Brownian motion.

We generate λ_t as follows. First, we draw a vector λ uniformly from $[0, 100]^I$. Depending on whether the experiment tests the no-forecast algorithm or not, we then either set $\lambda_{i,t} = M\lambda_i$ (for the no forecast case) or we set $\lambda_{i,t} = M\lambda_i(1 + 2t)$ with probability 1/2 and $\lambda_{i,t} = M\lambda_i(1 - t/2)$ with the remaining probability (in the case where we test the algorithm incorporating forecasts). In both cases, M is selected so that the load factor LF takes the appropriate value for the instance we wish to generate.

We set σ so that CV takes the appropriate value; we use bisection to find the appropriate value of σ here.

Finally, we generate the stochastic process via the natural discretization of the continuous time process; in particular, we use the recursion

$$\bar{\Lambda}_{n\Delta} = \lambda_{n\Delta} + (1 - \Delta)(\bar{\Lambda}_{(n-1)\Delta} - \lambda_{(n-1)\Delta}) + \sigma\epsilon_n,$$

where $\Delta = 1/100$ and ϵ_n is a zero mean normal with variance Δ .

5.2. Results for instances from generative family. We first consider problems generated from the family described with λ set to a constant. We consider 30 ensembles of instances. Each ensemble differs in the parameters (LF, CV) and itself contains 30 individual problem instances. We employ the reoptimization scheme designed for scenarios where no forecast is available, taking the reoptimization frequency, N , to be 100. We use as our upper bound the quantity $\mathbb{E}[J_{\{\lambda_t\}}^{\text{UB}}(x_0)]$.

The results of these experiments are summarized in Table H.1. Here we make several observations: First, performance relevant to the clairvoyant upper bound is consistently good; it is at least within 80% of this upper bound, and, frequently, well within 90%. We observe some performance degradation in regimes of extremely high volatility. Also, problems with low load factors (i.e., where demand is scarce) appear to be more challenging for the scheme. This is somewhat intuitive if seen from the perspective that in such a regime, one will not be able to consume each component of x_0 with its “optimal” impression.

Next, we consider a similar ensemble of problems, but with λ_t allowed to be time varying (in the manner described in the previous section). We employ the reoptimization scheme that utilizes forecasts. This scheme requires a tuning parameter β . We consider two sets of experiments; the first uses the “robust” choice of β ($\beta = 0.406$) identified via Theorem 4.1; the results for this scheme are described in Table H.2. We then allow the algorithm designer to choose β from the set $\{0, 0.2, \dots, 1.0\}$, and report performance for the best of these values in Table H.3. Again, we set $N = 100$ and use $\mathbb{E}[J_{\{\lambda_t\}}^{\text{UB}}(x_0)]$ for upper bound comparisons in all experiments. In addition, we note that we employ the “inventory sharing” improvement described following Theorem 4.1.

The results for these experiments are described in Tables H.2 and H.3, respectively. Although it is not displayed here, the optimal β in essentially all cases for Table H.3 where demand was volatile was $\beta = 0$; i.e., the simple reoptimization scheme that ignores forecasts altogether. We see qualitatively similar results to the case where λ is a constant. Moreover, when one allows the user to optimize β , performance is essentially as good as the constant λ case.

Lastly, we test the no-drift algorithm by varying N in the set $\{1, 2, 5, 10, 50, 100\}$, while keeping the load factor constant at one. As expected, the impact of discretization increases with volatility. However, performance is often satisfactory even for $N = 1, 2$ or 5. Increasing N from 1 to 5 does not seem to have a reliable effect, however increasing N from these lower ranges to 100 monotonically and markedly improves performance. The results of the experiments are summarized in Table H.4.

5.3. A real instance from an ad platform. We also consider experiments on a real instance of the ad display problem described here. In particular, we have data from a mobile ad platform for a single day of traffic. On this day, the platform served 240 distinct advertisers using impressions from a highly heterogeneous pool. Each campaign can only be served by a subset of this traffic based on a set of parameters. Payments are uniform across compatible traffic for a given advertiser; i.e., $p_{(i,a)} \in \{0, p_a\}$. We aggregated impressions based on their originating website/mobile application resulting in a total of about 40 traffic sources. The number of arrivals during the 24-hour interval varies from four million for the largest inventory type, to 10 thousand for the smallest. The coefficient of variation varies from 0.4 to 1.5 and there are noticeable intraday, cyclical trends in arrival rates (for example, arrivals have larger intensity in the morning and evening, versus late night). The average in-degree of the implied bipartite graph from inventory to campaigns is 4.5, the average out-degree is 18 and the load factor of user requests to campaign capacities is 1.5.

We chose to use the reoptimization scheme without forecast inputs and set the reoptimization frequency to $N = 24$ (i.e., once every hour). Solutions to the LP were interpreted probabilistically, and we used as our benchmark the clairvoyant bound (which in this case corresponds to assigning the day's traffic after it has been realized). Our scheme earns revenues that are within 99.3% of this benchmark, which is encouraging. Moreover, performance within 99% of the benchmark was maintained under two stress tests, namely, scaling down the load factor to 0.75, and introducing heterogeneous prices (where we multiplied prices by a standard uniform random variable).

6. Conclusion. The main contribution of this paper has been to develop a simple, easy to interpret algorithm that can efficiently solve a large class of dynamic allocation problems. Our method is robust (as witnessed by worst-case guarantees) and in the event that demand volatility (or, equivalently, deviations of demand from its forecast) is not large, the scheme is simultaneously optimal. Practical experiments have shown that the approach is promising both in terms of performance and practicality. At a somewhat more abstract level, we believe this work contributes to the (theoretically) poorly understood area of model predictive control, and as such, we believe the simple analytical tools developed here (the balancing property) may provide value in other contexts. There are many directions that merit further attention:

(i) *Practical scalability.* It would be interesting to examine heuristics that do not have to explicitly re-solve a new linear program at each opportunity to reoptimize. In particular, if the basis does not change dramatically, there should be some value to local search around the incumbent basis.

(ii) *Reoptimization frequency.* How large should N be? We established that $N = 1$ (i.e., no reoptimization), can be arbitrarily bad, and that N large yields good performance guarantees. In our experiments, relatively small values of N seemed viable. Ideally the choice of N should be driven by the volatility in the underlying demand process (i.e., σ in our experiments), and understanding this dependence would be interesting (but potentially quite nontrivial).

(iii) *Inventory balancing.* In addition to being a potentially valuable theoretical tool in other contexts, the inventory balancing property of our reoptimization scheme permits some interesting (unintended) applications. For instance, in the case of the ad display problem, it automatically results in allocations that satisfy what are commonly called “pacing” constraints where resources (eg., budgets) may be consumed at most at some prespecified uniform rate over time. It would be interesting to explore the applications of this property further.

Acknowledgments. The authors thank Steve Graves for his input in the early stages of this work. The authors acknowledge valuable technical input from Yiwei Chen. In addition, the second author acknowledges useful discussions with Hamid Nazerzadeh, Devavrat Shah, and David Yao.

Appendix A. Proof for §2

PROOF OF LEMMA 2.1. Employing the notation $\bar{\sigma}_{T,1} \triangleq (1/T) \int_0^T \sigma_t dt$, we have

$$\begin{aligned} \frac{E[(1/T) \int_0^T f(\Lambda_t) dt]}{f(E[(1/T) \int_0^T \Lambda_t dt])} &= \frac{E[(1/T) \int_0^T f(\Lambda_t) dt]}{f((1/T) \int_0^T E[\Lambda_t] dt)} \\ &\geq \frac{E[(1/T) \int_0^T f(\Lambda_t) dt]}{f((1/T) \int_0^T (\lambda_t + \sigma_t/\sqrt{2\pi}) dt)} \\ &= \frac{1}{T} \int_0^T \left[\int_{-\infty}^{\infty} \frac{f((\lambda + y)^+)}{f(\lambda + (1/T) \int_0^T (\sigma_t dt)/\sqrt{2\pi})} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy \right] dt \end{aligned}$$

$$\begin{aligned}
 &\geq \frac{1}{T} \int_0^T \left[\int_{-\infty}^{\infty} \min \left\{ \frac{(\lambda + y)^+}{\lambda + (1/T) \int_0^T (\sigma_t dt) / \sqrt{2\pi}}, 1 \right\} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy \right] dt \\
 &= \frac{1}{T} \int_0^T \left[1 - \Phi \left(\frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) + \int_{-\lambda}^{\bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{\lambda + y}{\lambda + \bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy \right] dt \\
 &\geq \frac{1}{T} \int_0^T \left[1 - \Phi \left(\frac{\bar{\sigma}_{T,1}}{\sigma_t \sqrt{2\pi}} \right) + \int_0^{\bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{y}{\lambda + \bar{\sigma}_{T,1}/\sqrt{2\pi}} \frac{\exp(-y^2/2\sigma_t^2)}{\sqrt{2\pi\sigma_t^2}} dy \right] dt \\
 &\geq 0.342.
 \end{aligned}$$

In this sequence of bounds, the first inequality comes from the fact that $f(\cdot)$ was assumed nondecreasing, and the fact that for a normal random variable with mean μ and variance σ^2 , one has

$$E[X^+] \leq \mu + \frac{\sigma}{\sqrt{2\pi}},$$

so that $E[\Lambda_t] \leq \lambda_t + \sigma_t/\sqrt{2\pi}$. The second equality follows from Fubini's theorem. The second inequality follows from property 3 of Lemma 3.3 applied to $f(\cdot)$. The final inequality is Lemma 14 of Chen and Farias [12]. \square

Appendix B. Proof for §3.2

PROOF OF PROPOSITION 3.1. Let $\{z_t\} \in \Pi_N$ be an ϵ -optimal admissible control policy feasible for (1) (which exists for arbitrary $\epsilon > 0$, see Bertsekas and Shreve [6]). For a given sample path ω , define

$$\hat{z}_e = \frac{\int_0^T \Lambda_{i(e),t} z_{e,d(t)} \mathbf{1}_{I_{e,t}} dt}{\int_0^T \Lambda_{i(e),t} dt}.$$

Then \hat{z}_e is feasible for (2) by the definition of $I_{e,t}$ so that it yields a solution to (2) of value no greater than $J_{\{\Lambda_t\}}^*(x_0)$. Moreover, the value of this solution is precisely the value garnered by $\{z_t\}$ on ω . Taking expectations yields $E[J_{\{\Lambda_t\}}^*(x_0)] \geq J^{*,N}(x_0) - \epsilon$. Because our choice of $\epsilon > 0$ was arbitrary, the result follows. \square

Appendix C. Proofs for §3.3

PROOF OF LEMMA 3.1. Observe that the quantity of resource type k consumed by edge e in the interval $[jT/N, (j+1)T/N)$ is at least

$$\frac{T}{N} \left(\Lambda_{i(e),jT/N}^{\min} z_{e,jT/N}^R A_{k,e} - \sum_{e' \neq e} (\Lambda_{i(e'),jT/N}^{\max} - \Lambda_{i(e'),jT/N}) z_{e',jT/N}^R A_{k,e'} \right).$$

It follows that the revenues garnered along edge e over that interval are at least

$$\frac{T}{N} p_e (\Lambda_{i(e),jT/N}^{\min} z_{e,jT/N}^R - C_e \sum_{e' \neq e} (\Lambda_{i(e'),jT/N}^{\max} - \Lambda_{i(e'),jT/N})),$$

where $C_e = \max_{\{e',k:A_{k,e'}>0\}} A_{k,e'}/A_{k,e}$. It then follows that total revenues over the interval are at least

$$\begin{aligned}
 &\frac{T}{N} \left(\sum_e p_e z_{e,jT/N}^R \Lambda_{i(e),jT/N}^{\min} + \sum_e p_e z_{e,jT/N}^R \Lambda_{i(e),jT/N}^{\min} - \sum_e p_e z_{e,jT/N}^R \Lambda_{i(e),jT/N} - \sum_e C_e p_e \sum_{e' \neq e} (\Lambda_{i(e'),jT/N}^{\max} - \Lambda_{i(e'),jT/N}) \right) \\
 &\geq \frac{T}{N} \left(\sum_e p_e z_{e,jT/N}^R \Lambda_{i(e),jT/N}^{\min} - C \left(\sum_e p_e \right) \sum_e (\Lambda_{i(e),jT/N}^{\max} - \Lambda_{i(e),jT/N}^{\min}) \right) \\
 &= \frac{T}{N} \left(\frac{\text{LP}(jT/N, \Lambda_{jT/N}, x_{jT/N})}{T - jT/N} - C \left(\sum_e p_e \right) \sum_e (\Lambda_{i(e),jT/N}^{\max} - \Lambda_{i(e),jT/N}^{\min}) \right),
 \end{aligned}$$

where $C = \max C_e$. Summing over intervals yields the result. \square

PROOF OF COROLLARY 3.1. Given the continuity of the sample path $\Lambda_{i,t}$ in t for all i , we have that

$$\lim_N \frac{T}{N} \sum_{j=0}^{N-1} \Lambda_{i,jT/N}^{\max} = \lim_N \frac{T}{N} \sum_{j=0}^{N-1} \Lambda_{i,jT/N}^{\min}$$

for all i . The result then follows from Lemma 3.1. \square

PROOF OF COROLLARY 3.2. We begin with making two elementary observations. First,

$$LP(\Lambda, t, x - \delta) \geq LP(\Lambda, t, x) - \sum_k \delta_k \frac{\sum_e p_e}{\min_{e: A_{k,e} > 0} A_{k,e}}.$$

Second, $LP(t, \Lambda, x)$ is a component-wise monotone function of x . These observations with the result of the balancing lemma then immediately yield

$$\begin{aligned} & \frac{LP(jT/N, \Lambda_{jT/N}, x_{jT/N})}{T - jT/N} \\ & \geq \frac{LP(jT/N, \Lambda_{jT/N}, x_0(N - j)/N)}{T - jT/N} + \sum_k \left[\sum_{l=0}^{j-1} \sum_e \frac{T}{N} A_{k,e} (\Lambda_{i(e),lT/N} - \Lambda_{i(e),lT/N}^{\max}) \right] \frac{\sum_e p_e}{\min_{e: A_{k,e} > 0} A_{k,e}} \\ & \geq \frac{LP(jT/N, \Lambda_{jT/N}, x_0(N - j)/N)}{T - jT/N} + M \left[\sum_{l=0}^{N-1} \sum_i \frac{T}{N} (\Lambda_{i,lT/N}^{\min} - \Lambda_{i,lT/N}^{\max}) \right], \end{aligned}$$

where $M \triangleq (\max_{k,e} A_{k,e} / \min_{k,e: A_{k,e} > 0} A_{k,e}) (\sum_e p_e) KE$. Consequently,

$$\begin{aligned} & \frac{T}{N} \sum_{j=0}^{N-1} \frac{LP(jT/N, \Lambda_{jT/N}, x_{jT/N})}{T - jT/N} \\ & \geq \frac{T}{N} \sum_{j=0}^{N-1} \left(\frac{LP(jT/N, \Lambda_{jT/N}, x_0(N - j)/N)}{T - jT/N} + M \left[\sum_{l=0}^{N-1} \sum_i \frac{T}{N} (\Lambda_{i,lT/N}^{\min} - \Lambda_{i,lT/N}^{\max}) \right] \right) \\ & = \frac{T}{N} \left(\sum_{j=0}^{N-1} \frac{LP(jT/N, \Lambda_{jT/N}, x_0(N - j)/N)}{T - jT/N} + MT \sum_{l=0}^{N-1} \sum_i (\Lambda_{i,lT/N}^{\min} - \Lambda_{i,lT/N}^{\max}) \right). \end{aligned} \tag{C1}$$

Then, using the fact that the continuity of the sample path $\Lambda_{i,t}$ in t for all i yields

$$\lim_N \frac{T}{N} \sum_{j=0}^{N-1} \Lambda_{i,jT/N}^{\max} = \lim_N \frac{T}{N} \sum_{j=0}^{N-1} \Lambda_{i,jT/N}^{\min}$$

for all i , we may take the limit infimum on both sides of (C1) to arrive at the result. \square

Appendix D. Proofs for §3.4

PROOF OF LEMMA 3.3. We have

- (i) that $f(0) = 0$ and f is continuous and nondecreasing follows immediately from the definition of f ;
- (ii) $f(\cdot)$ is concave because $\min\{w/\Lambda, 1\}$ is concave, and because summations preserve concavity;
- (iii) for $w \geq v$, $\min\{w/v, 1\} = 1$ and $f(w)/f(v) \geq 1$ by virtue of f being nondecreasing. For $w < v$, because f is concave, and $w, v > 0$ with $w/v < 1$ we have $f(w) = f(0 + w/v \cdot v) \geq w/v \cdot f(v)$. Thus, $f(w)/f(v) \geq w/v$.
- (iv) The result is a consequence of Jensen’s inequality. \square

PROOF OF LEMMA 3.4. Define $\tilde{z} \in \mathbb{R}^E$ according to $\tilde{z}_e = z_e(t, \Lambda, x) \min\{\Lambda_{i(e)}/u_{i(e)}, 1\}$. Observe that \tilde{z} is a feasible solution to $LP(t, u, x)$. In particular,

$$\begin{aligned} \sum_e A_{k,e} \tilde{z}_e u_{i(e)} (T - t) &= \sum_e A_{k,e} z_e(t, \Lambda, x) \min\{\Lambda_{i(e)}/u_{i(e)}, 1\} u_{i(e)} (T - t) \\ &\leq \sum_e A_{k,e} z_e(t, \Lambda, x) \Lambda_{i(e)} (T - t) \\ &\leq x_k, \end{aligned}$$

where the last inequality is due to the feasibility of $z(t, \Lambda, x)$ for $LP(t, \Lambda, x)$. Further, $\tilde{z} \in \mathcal{Z}$ by construction. Now, the objective value of this solution is precisely $\sum_{i=0}^I f_i^{(t, \Lambda, x)}(u_i)$, and the result follows because the value of an optimal solution to $LP(t, u, x)$ should be at least as large. \square

Appendix E. Proofs for §3.6

PROOF OF PROPOSITION 3.2. Begin by considering a policy $\{z_t^U\}$ defined according to $z_{t,(1,1)}^U = 0, z_{t,(2,1)}^U = 1$. Clearly, $J^U(x_0) = T^{4/3+\epsilon}/2$. Consequently,

$$J^{UB}(x_0) \geq T^{4/3+\epsilon}/2. \tag{E1}$$

Now consider our reoptimization policy with $N = 1$. It is easy to see that for this policy, we have $z_{t,(1,1)}^R = z_{t,(2,1)}^R = 1$. We now compute an upper bound on $J^{R,1}(x_0)$. First, note that $E[\Lambda_{1,s}] = \frac{1}{2}\sqrt{s}$ and $\text{Var}[\Lambda_{1,s}] = s(\pi - 1)$. Hence, for any $0 \leq t \leq T$, the quantity $\int_0^t \Lambda_{1,s} ds$ has moments $E[\int_0^t \Lambda_{1,s} ds] = t/2 + \frac{3}{2}t^{3/2}$ and $\text{Var}[\int_0^t \Lambda_{1,s} ds] = (\pi - 1/2)t^2$. By Chebyshev's inequality, it follows that

$$P\left[\left|\int_0^t \Lambda_{1,s} ds - \left(\frac{t}{2} + \frac{3}{2}t^{3/2}\right)\right| \geq \frac{3}{4}t^{3/2}\right] \leq \frac{8(\pi - 1)}{9} \frac{1}{t}.$$

In particular, with probability at least $1 - 8(\pi - 1)/(9t)$, $\int_0^t \Lambda_{1,s} ds \geq t/2 + (3/4)t^{3/2}$. Now, define τ to be the solution to the equation $\tau + (3/4)\tau^{3/2} = T$. One may consequently interpret τ as the first time t at which $x_t = 0$ assuming the Λ_1 process followed its mean.

Now define the stopping time $\hat{\tau}$ according to $\hat{\tau} = \inf\{t: \int_0^t \Lambda_{1,s} + \Lambda_{2,s} ds \geq T\}$. Now on the event where $\hat{\tau} \leq \tau$, we have

$$\begin{aligned} \int_0^{\hat{\tau}} \Lambda_{1,s} ds &= T - \int_0^{\hat{\tau}} \Lambda_{2,s} ds \\ &\geq T - \int_0^{\tau} \Lambda_{2,s} ds \\ &= \frac{\tau}{2} + \frac{3}{4}\tau^{3/2}. \end{aligned}$$

Consequently,

$$\begin{aligned} \int_0^{\hat{\tau}} \Lambda_{2,s} ds &= T - \int_0^{\hat{\tau}} \Lambda_{1,s} ds \\ &\leq T - \frac{\tau}{2} - \frac{3}{4}\tau^{3/2} \\ &= \frac{\tau}{2}. \end{aligned}$$

It follows that on the event where $\tau \leq \hat{\tau}$, we have

$$\begin{aligned} J_{\{\Lambda_t\}}^{R,1}(x_0) &= T^{1/3+\epsilon} \int_0^{\hat{\tau}} \Lambda_{2,s} + (T - \int_0^{\hat{\tau}} \Lambda_{2,s} ds) \\ &\leq T^{1/3+\epsilon} \frac{\tau}{2} + T - \frac{\tau}{2}. \end{aligned}$$

On the complementary event, we consider the trivial upper bound

$$J_{\{\Lambda_t\}}^{R,1}(x_0) \leq T^{4/3+\epsilon}.$$

Now notice that $P(\hat{\tau} \leq \tau) = P(\int_0^{\tau} \Lambda_{1,s} + \Lambda_{2,s} ds \geq T) \geq 1 - 8(\pi - 1)/(9\tau)$, so that taking expectations immediately yields with the two inequalities:

$$J^{R,1}(x_0) \leq \left(1 - \frac{8(\pi - 1)}{9\tau}\right) \left(T^{1/3+\epsilon} \frac{\tau}{2} + T - \frac{\tau}{2}\right) + \frac{8(\pi - 1)}{9\tau} T^{4/3+\epsilon}. \tag{E2}$$

But it is easy to verify that $\tau = \Theta(T^{2/3})$, so that the right-hand side of the inequality is in fact $\Theta(T^{1+\epsilon})$. The result then follows immediately from (E2) and (E1). \square

PROOF OF THEOREM 3.6. First, let us define

$$f(N) = \frac{1}{N} \sum_{j=0}^{N-1} \sum_{i=1}^J (\Lambda_{i,jT/N}^{\max} - \Lambda_{i,jT/N}^{\min}).$$

We begin with a simple proposition

LEMMA E1.

$$\frac{T}{N} \sum_{j=0}^{N-1} \frac{\text{LP}(0, \Lambda_{jT/N}, x_0)}{T} - \frac{1}{T} \int_0^T \text{LP}(0, \Lambda_t, x_0) dt \geq -\max_e p_e f(N).$$

PROOF. We have

$$\begin{aligned} \frac{T}{N} \sum_{j=0}^{N-1} \frac{\text{LP}(0, \Lambda_{jT/N}, x_0)}{T} - \frac{1}{T} \int_0^T \text{LP}(0, \Lambda_t, x_0) dt &\geq \frac{1}{T} \sum_{j=0}^{N-1} \left(\text{LP}(0, \Lambda_{jT/N}, x_0) \frac{T}{N} - \text{LP}(0, \Lambda_{jT/N}^{\max}, x_0) \frac{T}{N} \right), \\ &\geq -\max_e p_e f(N) \end{aligned}$$

where the first inequality follows from the monotonicity of $\text{LP}(\cdot, \cdot, \cdot)$ in its second argument, and the final inequality follows from the fact that for any $\delta \geq 0$ and any $u, x \geq 0$, we must have

$$\text{LP}(0, u + \delta, x) - \text{LP}(0, u, x) \leq \max_e p_e \sum_i \delta_i. \quad \square$$

We next define a few useful “error terms,” namely, $h_1(N) = TC(\sum_e p_e)Ef(N)$, $h_2(N) = MT^2f(N)$, and finally $h_3(N) = \max_e p_e f(N)$. The constants C and M are defined in Lemma 3.1 and Corollary 3.2, respectively. It will be useful to characterize the rate at which these terms approach zero. In particular, define

$$g(N) = h_1(N) + h_2(N) + h_3(N).$$

We have

LEMMA E2.

$$\limsup_N \frac{g(N)}{\sqrt{2 \log N/N}} \leq C' \sum_i \sigma_i \quad \text{a.s.},$$

where $C' \triangleq TC(\sum_e p_e)E + MT^2 + \max_e p_e$. Consequently,

$$\limsup_N \frac{E[g(N)]}{\sqrt{2 \log N/N}} \leq C' \sum_i \sigma_i.$$

PROOF. We have that for almost all ω and N sufficiently large,

$$\begin{aligned} g(N) &= C' f(N) \\ &\leq C' \sum_i \frac{1}{N} \sum_{j=0}^{N-1} \max_j (\Lambda_{i,jT/N}^{\max} - \Lambda_{i,jT/N}^{\min}) \\ &= C' \sum_i \max_j (\Lambda_{i,jT/N}^{\max} - \Lambda_{i,jT/N}^{\min}). \end{aligned}$$

From the modulus of continuity of the sample paths of the rate process, Theorem 3.5, we have that

$$\limsup_N \frac{g(N)}{\sqrt{2 \log N/N}} \leq C' \sum_i \sigma_i \quad \text{a.s.}$$

An application of the reverse Fatou’s Lemma yields

$$\limsup_N \frac{E[g(N)]}{\sqrt{2 \log N/N}} \leq C' \sum_i \sigma_i. \quad \square$$

We are now in a position to prove our result. In particular

$$\begin{aligned} J_{\{\Lambda_i\}}^{\text{R},N}(x_0) &\geq \frac{T}{N} \sum_{j=0}^{N-1} \frac{\text{LP}(jT/N, \Lambda_{jT/N}, x_{jT/N})}{T - jT/N} - h_1(N) \\ &\geq \frac{T}{N} \sum_{j=0}^{N-1} \frac{\text{LP}(jT/N, \Lambda_{jT/N}, x_0(N-j)/N)}{T - jT/N} - h_1(N) - h_2(N) \\ &= \frac{T}{N} \sum_{j=0}^{N-1} \frac{\text{LP}(0, \Lambda_{jT/N}, x_0)}{T} - h_1(N) - h_2(N) \\ &\geq \frac{1}{T} \int_0^T \text{LP}(0, \Lambda_t, x_0) dt - h_1(N) - h_2(N) - h_3(N), \end{aligned}$$

where the first inequality follows from the proof of Lemma 3.1, the second inequality follows from the proof of Corollary 3.2, the equality follows from the fact that $\text{LP}(t, \Lambda, x(1-t/T)) = ((T-t)/T)\text{LP}(0, \Lambda, x)$, and the

final inequality follows from Lemma E1. Now, taking expectations yields

$$\begin{aligned} \mathbb{E}[J_{\{\Lambda_t\}}^{\mathbb{R},N}(x_0)] &\geq \mathbb{E}\left[\frac{1}{T} \int_0^T \text{LP}(0, \Lambda_T, x_0) dt\right] - \mathbb{E}[g(N)] \\ &\geq 0.342 \mathbb{E}[J_{\{\Lambda_t\}}^{\text{UB}}(x_0)] - \mathbb{E}[g(N)], \end{aligned}$$

where the second inequality follows from the proof of our approximation guarantee in the setting with a large number of reoptimizations, Theorem 3.2. Defining $\Delta(N) = \mathbb{E}[g(N)]$ with the result of Lemma E2 yields the proof of the theorem.

Appendix F. Proof for §4.2

PROOF OF LEMMA 4.1. Let us define $\bar{p} \triangleq \max_e p_e$ and $\bar{A} \triangleq \max_{k,e} A_{k,e}$. We begin by observing that

$$\begin{aligned} J_{\{\Lambda_t\}}^{\mathbb{R}-\beta}(x_0) &= \int_0^T \sum_e p_e ((1-\beta)\hat{z}_{e,d(t)}^{\mathbb{R}} + \beta\hat{z}_{e,d(t)}^{\mathbb{D}}) \Lambda_{i(e),t} \mathbb{1}_{\{I_e,t\}} dt \\ &\geq \int_0^T \sum_e p_e ((1-\beta)\hat{z}_{e,d(t)}^{\mathbb{R}} + \beta\hat{z}_{e,d(t)}^{\mathbb{D}}) \Lambda_{i(e),t} dt - \bar{p} \sum_k \left(\sum_e \int_0^T A_{k,e} \Lambda_{i(e),t} \hat{z}_{e,t}^{\mathbb{R}-\beta} dt - x_{0,k} \right)^+ \\ &\geq (1-\beta)J_{S,\{\Lambda_t\}}^{\mathbb{R}}(x_0) + \beta J_{D,\{\Lambda_t\}}^{\mathbb{R}}(x_0) - \bar{p} \sum_k \left(\sum_e \int_0^T A_{k,e} (1-\beta) \Lambda_{i(e),t} \hat{z}_{e,t}^{\mathbb{R}} dt - (1-\beta)x_{0,k} \right)^+ \\ &\quad - \bar{p} \sum_k \left(\sum_e \int_0^T A_{k,e} \beta \Lambda_{i(e),t} \hat{z}_{e,t}^{\mathbb{D}} dt - \beta x_{0,k} \right)^+. \end{aligned} \tag{F1}$$

Next, observe that

$$\begin{aligned} \sum_e \int_0^T A_{k,e} \Lambda_{i(e),t} \hat{z}_{e,t}^{\mathbb{D}} dt &\leq \frac{T}{N} \sum_{j=0}^{N-1} \sum_e \hat{z}_{e,jT/N}^{\mathbb{D}} A_{k,e} \Lambda_{i(e),jT/N}^{\max} \\ &\leq \frac{T}{N} \left(\sum_{j=0}^{N-1} \sum_e \hat{z}_{e,jT/N}^{\mathbb{D}} A_{k,e} \Lambda_{i(e),jT/N} + \bar{A} \sum_{j=0}^{N-1} \sum_e \Lambda_{i(e),jT/N}^{\max} - \Lambda_{i(e),jT/N} \right) \\ &\leq \frac{T}{N} \left(\sum_{j=0}^{N-1} \sum_e \hat{z}_{e,jT/N}^{\mathbb{D}} A_{k,e} \Lambda_{i(e),jT/N} + \bar{A} \sum_{j=0}^{N-1} \sum_e \Lambda_{i(e),jT/N}^{\max} - \Lambda_{i(e),jT/N} \right). \end{aligned} \tag{F2}$$

But for any i ,

$$\lim_N \frac{T}{N} \sum_{j=0}^{N-1} \Lambda_{i,jT/N}^{\max} - \Lambda_{i,jT/N} = 0$$

by virtue of the continuity of $\Lambda_{i,t}$ and further,

$$\lim_N \frac{T}{N} \sum_{j=0}^{N-1} \hat{z}_{e,jT/N}^{\mathbb{D}} A_{k,e} \Lambda_{i(e),jT/N} = z_e^{\mathbb{D}} A_{k,e} \int_0^T \lambda_{i(e),t} dt,$$

since $\lambda_{i,t}$ is continuous in t for all i . Since, by the definition of $z^{\mathbb{D}}$, we have $\sum_e z_e^{\mathbb{D}} A_{k,e} \int_0^T \lambda_{i(e),t} dt \leq x_{0,k}$, it then follows from (F2) that

$$\limsup_N \sum_e \int_0^T A_{k,e} \Lambda_{i(e),t} \hat{z}_{e,t}^{\mathbb{D}} dt \leq x_{0,k}$$

for all k . Consequently, for all k ,

$$\lim_N \left(\sum_e \int_0^T A_{k,e} \beta \Lambda_{i(e),t} \hat{z}_{e,t}^{\mathbb{D}} dt - \beta x_{0,k} \right)^+ = 0.$$

Now, since

$$\left(\sum_e \int_0^T A_{k,e} \beta \Lambda_{i(e),t} \hat{z}_{e,t}^{\mathbb{D}} dt - \beta x_{0,k} \right)^+ \leq E \bar{A} \sum_i \max_{t \in [0,T]} \Lambda_{i,t}$$

the dominated convergence theorem allows us to conclude that for all k ,

$$\lim_N \mathbb{E} \left[\left(\sum_e \int_0^T A_{k,e} \beta \Lambda_{i(e),t} \hat{z}_{e,t}^{\mathbb{D}} dt - \beta x_{0,k} \right)^+ \right] = 0. \tag{F3}$$

Further, we must have that

$$\sum_e \int_0^T A_{k,e} \Lambda_{i(e),t} \hat{z}_{e,t}^R dt \leq x_{0,k} + \max_{j,i} \frac{T}{N} (\Lambda_{i,jT/N}^{\max} - \Lambda_{i,jT/N}^{\min}) \bar{A}$$

by virtue of the dynamics of \hat{x}_t^R and the definition of $\hat{z}_{jT/N}^R$, which guarantees that $A_{k,e} \hat{z}_{e,d(t)}^R = 0$ if $\hat{x}_{k,d(t)}^R = 0$. It follows that

$$\lim_N \left(\sum_e \int_0^T A_{k,e} (1 - \beta) \Lambda_{i(e),t} \hat{z}_{e,t}^R dt - (1 - \beta) x_{0,k} \right)^+ = 0.$$

As in (F3), the dominated convergence theorem applies to yield

$$\lim_N \mathbb{E} \left[\left(\sum_e \int_0^T A_{k,e} (1 - \beta) \Lambda_{i(e),t} \hat{z}_{e,t}^R dt - (1 - \beta) x_{0,k} \right)^+ \right] = 0. \tag{F4}$$

Taking an expectation followed by the limit infimum on both sides of (F1) then yields by (F3) and (F4) the result. \square

Appendix G. Proofs for §4.3

PROOF OF LEMMA 4.2. Let us define the process $\tilde{\Lambda}_t$ according to $\tilde{\Lambda}_t = (\Lambda_t - \lambda_t)^+$ and observe that $\tilde{\Lambda}_t$ satisfies the conditions of Theorem 3.2. Now, we have

$$\begin{aligned} \liminf_N \frac{\mathbb{E}[J_{S, \{\Lambda_t\}}^{R-\beta}(x_0)]}{\mathbb{E}[J_{\{\tilde{\Lambda}_t\}}^{UB}(x_0)]} &\geq \frac{\mathbb{E}[\liminf_N (T/N) \sum_{j=0}^{N-1} ((LP(jT/N, \Lambda_{jT/N}, \hat{x}_0^R(N-j)/N)) / (T - jT/N))]}{\mathbb{E}[J_{\{\tilde{\Lambda}_t\}}^{UB}(x_0)]} \\ &\geq \frac{\mathbb{E}[\liminf_N (T/N) \sum_{j=0}^{N-1} ((LP(jT/N, \tilde{\Lambda}_{jT/N}, \hat{x}_0^R(N-j)/N)) / (T - jT/N))]}{\mathbb{E}[J_{\{\tilde{\Lambda}_t\}}^{UB}(x_0)]}, \end{aligned}$$

where the first inequality is a consequence of Fatou’s lemma and (6), and the second inequality follows from the monotonicity of $LP(t, \Lambda, x)$ in Λ , and the fact that $\Lambda_t \geq \tilde{\Lambda}_t$. The remainder of the proof is then identical to Theorem 3.2. \square

PROOF OF LEMMA 4.4. We have

$$\begin{aligned} J_{\{\lambda_t\}}^{UB}(x_0) + J_{\{(\tilde{\Lambda}_t - \lambda_t)^+\}}^{UB}(x_0) &= LP\left(0, \frac{1}{T} \int_0^T \lambda_t dt, x_0\right) + LP\left(0, \frac{1}{T} \int_0^T (\tilde{\Lambda}_t - \lambda_t)^+ dt, x_0\right) \\ &\geq LP\left(0, \frac{1}{T} \int_0^T \lambda_t + (\tilde{\Lambda}_t - \lambda_t)^+ dt, x_0\right) \end{aligned}$$

Appendix H. Tables of experimental results We summarize our experimental results in the following tables. The 95% confidence intervals for reported figures are within $\pm 5\%$.

TABLE H.1. No-drift algorithm performance vs. upper bound.

Load Factor/CV	0 (%)	0.5 (%)	1 (%)	2.5 (%)	5 (%)	10 (%)
0.1	100.00	99.97	99.00	92.04	79.06	81.00
0.5	100.00	98.75	97.03	90.32	82.31	82.63
1	100.00	99.15	96.86	91.08	86.51	84.54
2	100.00	99.91	99.54	97.34	91.98	87.88
5	100.00	99.92	99.68	98.85	97.23	94.26

Note. The load factor is along the vertical axis and the CV is along the horizontal axis.

TABLE H.2. Drift algorithm performance vs. upper bound for $\beta = 0.406$.

Load Factor/CV	0 (%)	0.5 (%)	1 (%)	2.5 (%)	5 (%)	10 (%)
0.1	100.00	74.45	67.80	62.79	62.54	74.35
0.5	99.91	91.66	84.59	73.45	69.35	74.80
1	99.74	95.46	91.03	81.54	76.95	78.08
2	99.86	97.63	95.44	91.13	86.34	84.37
5	99.95	99.06	98.46	97.23	95.15	92.43

Note. The load factor is along the vertical axis and the CV is along the horizontal axis.

TABLE H.3. Drift algorithm performance vs. upper bound for optimized β .

Load Factor/CV	0 (%)	0.5 (%)	1 (%)	2.5 (%)	5 (%)	10 (%)
0.1	100.00	99.95	98.90	92.18	79.17	80.98
0.5	100.00	99.31	97.87	90.80	82.34	82.55
1	100.00	98.95	97.23	90.70	85.81	84.28
2	100.00	99.17	97.84	94.60	90.10	87.21
5	100.00	99.80	99.56	98.71	96.57	93.22

Note. The load factor is along the vertical axis and the CV is along the horizontal axis.

TABLE H.4. No-drift algorithm performance vs. upper bound for various re-solving frequencies.

Load Factor/CV	0 (%)	0.5 (%)	1 (%)	2.5 (%)	5 (%)	10 (%)
1	100.00	99.99	98.59	94.49	83.95	74.46
2	100.00	99.94	98.09	94.09	83.37	74.79
5	100.00	99.96	98.03	93.71	83.54	75.96
10	100.00	99.99	98.39	94.74	85.31	78.91
50	100.00	99.99	98.85	96.13	89.5	85.15
100	100.00	99.99	99.15	96.86	91.08	86.51

Note. The number of reoptimization intervals is along the vertical axis and the CV is along the horizontal axis.

$$\begin{aligned} &\geq \text{LP}\left(0, \frac{1}{T} \int_0^T \Lambda_t dt, x_0\right) \\ &= J_{\{\Lambda_t\}}^{\text{UB}}(x_0). \end{aligned}$$

The inequality follows from the concavity of $\text{LP}(t, u, x)$ in u and because any concave function $h: \mathbb{R}_+^n \rightarrow \mathbb{R}$ with $h(0) = 0$ must be super additive. The second inequality follows from the monotonicity of $\text{LP}(t, u, x)$ in u and since $\lambda_t + (\bar{\Lambda}_t - \lambda_t)^+ \geq \bar{\Lambda}_t$. Taking expectations in the inequality yields the result. \square

References

- [1] Adelman D (2007) Dynamic bid price in revenue management. *Oper. Res.* 55(4):647–661.
- [2] Agrawal S, Wang Z, Ye Y (2009) A dynamic near-optimal algorithm for online linear programming. *CoRR* abs/0911.2974.
- [3] Akan M, Ata B (2009) Bid-price controls for network revenue management: Martingale characterization of optimal bid prices. *Math. Oper. Res.* 34(4):912–936.
- [4] Ball MO, Queyranne M (2009) Toward robust revenue management: Competitive analysis of online booking. *Oper. Res.* 57(4):950–963.
- [5] Bemporad A (2006) Model-based predictive control design: New trends and tools. *CDC06: 45th Conf. Decision and Control, San Diego* (IEEE, New York), 6678–6683.
- [6] Bertsekas DP, Shreve SE (2007) *Stochastic Optimal Control: The Discrete-Time Case* (Athena Scientific, Nashua, NH).
- [7] Bertsimas D, Demir R (2002) An approximate dynamic programming approach to multidimensional knapsack problems. *Management Sci.* 48:550–565.
- [8] Besbes O, Maglaras C (2009) Revenue optimization for a make-to-order queue in an uncertain market environment. *Oper. Res.* 57(6):1438–1450.
- [9] Bramson M (1998) State space collapse with application to heavy traffic limits for multiclass queueing networks. *Queueing Systems* 30:89–148.
- [10] Buchbinder N, Naor J (2009) Online primal-dual algorithms for covering and packing. *Math. Oper. Res.* 34:270–286.
- [11] Chen H, Yao DD (1993) Dynamic scheduling of a multiclass fluid network. *Oper. Res.* 41:1104–1115.
- [12] Chen Y, Farias VF (2010) *Simple Policies for Dynamic Pricing with Imperfect Forecasts*. Preprint.
- [13] Farias VF, Van Roy B (2007) *An Approximate Dynamic Programming Approach to Network Revenue Management*. Preprint.
- [14] Farias VF, Van Roy B (2010) Dynamic pricing with a prior on market response. *Oper. Res.* 58(1):16–29.
- [15] Gallego G, van Ryzin G (1997) A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.* 45(1):24–41.
- [16] Jasin S, Kumar S (2012) A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Math. Oper. Res.* 37(2):313–345.
- [17] Karatzas I, Shreve SE (1991) *Brownian Motion and Stochastic Calculus* (Springer, New York).
- [18] Kleinberg R (2005) A multiple-choice secretary algorithm with applications to online auctions. *Proc. Sixteenth Annual ACM-SIAM Sympos. Discrete Algorithms. SODA '05* (Society for Industrial and Applied Mathematics, Philadelphia), 630–631.
- [19] Maglaras C, Meissner J (2006) Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing Service Oper. Management* 8(2):136–148.
- [20] Mahdian M, Nazerzadeh H, Saberi A (2012) Online optimization with uncertain information. *ACM Trans. Algorithms* 8(1):2.
- [21] Mehta A, Saberi A, Vazirani U, Vazirani V (2005) Adwords and generalized on-line matching. *FOCS'05: Proc. 46th Annual IEEE Sympos. Foundations Comput. Sci.* (IEEE Computer Society Press, Los Alamitos, CA), 264–273.
- [22] Reiman MI, Wang Q (2008) An asymptotically optimal policy for a quantity-based network revenue management problem. *Math. Oper. Res.* 33(2):257–282.
- [23] Shah D, Wischik D (2011) Switched networks with maximum weight policies: Fluid approximation and multiplicative state space collapse. *Ann. Appl. Probab.* Forthcoming.
- [24] Zhang D, Adelman D (2009) An approximate dynamic programming approach to network revenue management with customer choice. *Transportation Sci.* 43(3):381–394.