

Technical Note – Revenue Management with Calendar-Aware and Dependent Demands: Asymptotically Tight Fluid Approximations

Weiyan Li¹, Paat Rusmevichientong², Huseyin Topaloglu¹

¹School of Operations Research and Information Engineering, Cornell Tech, New York, NY 10044

²Marshall School of Business, University of Southern California, Los Angeles, CA 90089
wl425@cornell.edu, rusmevic@marshall.usc.edu, topaloglu@orie.cornell.edu

October 19, 2024

When modeling the demand in revenue management systems, a natural approach is to focus on a canonical interval of time, such as a week, so that we forecast the demand over each week in the selling horizon. Ideally, we would like to use random variables with general distributions to model the demand over each week. The current demand can give a signal for the future demand, so we also would like to capture the dependence between the demands over different weeks. Prevalent demand models in the literature, which are based on a discrete-time approximation to a Poisson process, are not compatible with these needs. In this paper, we focus on revenue management models that are compatible with a natural approach for forecasting the demand. Building such models through dynamic programming is not difficult. We divide the selling horizon into multiple stages, each stage being a canonical interval of time on the calendar. We have random number of customer arrivals in each stage, whose distribution is arbitrary and depends on the number of arrivals in the previous stage. The question we seek to answer is the form of the corresponding fluid approximation. We give the correct fluid approximation in the sense that it yields asymptotically optimal policies. The form of our fluid approximation is surprising as its constraints use expected capacity consumption of a resource up to a certain time period, conditional on the demand in the stage just before the time period in question. Letting K be the number of stages in the selling horizon, c_{\min} be the smallest resource capacity and ϵ be a lower bound on the mass function of the demand in a stage, we use the fluid approximation to give a policy with a performance guarantee of $1 - O\left(\frac{\sqrt{(c_{\min} + K/\epsilon^6) \log c_{\min}}}{c_{\min}}\right)$. As the resource capacities and number of stages increase with the same rate, the performance guarantee converges to one. To our knowledge, this result gives the first asymptotically optimal policy under dependent demands with arbitrary distributions. When the demands in different stages are independent, letting σ^2 be the variance proxy for the demand in each stage, a similar performance guarantee holds by replacing $\frac{1}{\epsilon^6}$ with σ^2 . Our computational work indicates that using the right fluid approximation can make a dramatic impact in practice.

Keywords: Fluid approximation, network revenue management, concentration inequality.

1. Introduction

A natural approach for modeling the demand in revenue management systems is to focus on a canonical interval of time, such as a week, so that we forecast the demand over each week in the selling horizon. Ideally, we would like to model the demand over each week through a random variable with an arbitrary distribution. Indeed, it is common for revenue management systems to produce forecasts stating that the demand, for example, during the week of July 17-23, 2023 has the log-normal distribution with mean 500 and standard deviation 250. These forecasts involve arbitrary demand distributions. Also, these forecasts are aware of the calendar in the sense that they have a concept of when each week ends and the next one starts. Furthermore, the current demand

usually gives a signal for the future demand, so we also would like to capture the dependence between the demands over different weeks. Prevalent demand models in the literature are not compatible with such a natural approach for forecasting the demand. In particular, a common demand model is based on dividing the selling horizon into a number of time periods such that there is at most one customer arrival at each time period and the arrivals at successive time periods are independent. This model corresponds to a discrete-time approximation to a Poisson process, but it has shortcomings. Under this model, the demand over an interval of time will always be approximately Poisson. To make matters worse, using λ_t to denote the probability that we have a customer arrival at time period t , over an interval of T time periods, the mean and standard deviation of the number of customer arrivals are, respectively, $\sum_{t=1}^T \lambda_t$ and $\sqrt{\sum_{t=1}^T \lambda_t (1 - \lambda_t)}$. Noting that $\sqrt{\sum_{t=1}^T \lambda_t (1 - \lambda_t)} \leq \sqrt{\sum_{t=1}^T \lambda_t}$, the ratio between the standard deviation and mean of the demand is at most $1/\sqrt{\sum_{t=1}^T \lambda_t}$. Thus, as the mean demand gets large, the coefficient of variation of the demand gets smaller. In other words, large demand variability and large demand volume cannot co-exist in this demand model. Finally, because this demand model is based on a Poisson process, the demands over different time intervals must be independent.

Motivated by the shortcomings discussed in the previous paragraph, we focus on revenue management models that are intrinsically compatible with a natural approach for forecasting the demand. Such a natural approach for forecasting the demand may specify the distribution of the demand over, for example, different weeks in the selling horizon, possibly along with the correlation structure between the demands in successive weeks. It is not too difficult to build these revenue management models through dynamic programming. We can divide the selling horizon into a number of stages, each stage representing a canonical interval of time on the calendar, such as a week. The number of customer arrivals in each stage is a random variable, whose distribution is arbitrary and depends on the number of customer arrivals in the previous stage. Therefore, along with the remaining capacities of the resources, the state variable in the dynamic program needs to keep track of the number of customer arrivals so far in the current stage and the number of customer arrivals in the previous stage, so that as a function of these two quantities, we can compute the probability of having one more customer arrival in the current stage. This dynamic program would give a precise specification of our model, but it is not useful for computing the optimal policy in practice because it involves a high-dimensional state variable.

We focus on fluid approximations for our model so that we can construct practical policies with performance guarantees and compute upper bounds on the optimal total expected revenues.

Our Results and Contributions: We start with a revenue management model based on a dynamic programming formulation that can handle arbitrary distributions for the number of

customer arrivals in each stage and allow dependence between the number of customer arrivals in successive stages. A stage may correspond to a canonical interval of time over which forecasts are produced. We capture the dependence between the demands through a Markov chain that specifies the distribution of the number of customer arrivals in one stage as a function of the number of customer arrivals in the previous stage. We give the correct fluid approximation for our revenue management model in the sense that the fluid approximation satisfies two properties. First, given that it is computationally difficult to find the optimal policy through a dynamic programming formulation, we can use our fluid approximation to construct approximate policies with performance guarantees. Second, we can use our fluid approximation to obtain an upper bound on the optimal total expected revenue, so that we can compare the total expected revenue of a heuristic policy with the upper bound to assess the optimality gap of the heuristic policy.

Structure of the Fluid Approximation. The structure of our fluid approximation turns out to be novel. In our fluid approximation, we use decision variables that allow the probability of accepting a customer request at a time period to depend on the number of customer arrivals in the previous stage. Due to the dependence between the demands in successive stages, this form of the decision variables is perhaps not surprising, but we are not aware of other fluid approximations with similar decision variables. More importantly, the constraints in our fluid approximation keep track of the expected capacity consumption of a resource up to a certain time period in a particular stage, conditional on the demand in the stage just before the time period in question. The conditional form of these constraints is surprising and does not appear in the literature. We show that the optimal objective value of our fluid approximation is an upper bound on the optimal total expected revenue. Thus, we can use this upper bound to assess the optimality gaps of heuristic policies.

Policies under Dependent Demands. Using our fluid approximation, we give an approximate policy. Letting K be the number of stages in the selling horizon, c_{\min} be the smallest resource capacity, L be the maximum number of resources used by a product and ϵ be a lower bound on the mass function of the demand in a stage given the demand in the previous stage, we show that our approximate policy has a performance guarantee of $\max \left\{ \frac{1}{4L}, \left(1 - 4 \frac{\sqrt{(c_{\min} + (K-1)/\epsilon^6) \log c_{\min}}}{c_{\min}} - \frac{L}{c_{\min}} \right) \right\}$. In many applications, the number of resources used by a particular product is usually small. In the airline setting, for example, the number of flight legs in an itinerary rarely exceeds two, corresponding to $L = 2$. By the first term in the max operator, our approximate policy has a constant-factor performance guarantee when L is uniformly bounded. By the second term in the max operator, as the number of stages and the capacities of the resources both increase linearly with rate θ , our performance guarantee converges to one with rate $1 - \frac{1}{\sqrt{\theta}}$.

Thus, our approximate policy is asymptotically optimal for systems that command large demands for the products and involve large capacities for the resources. During the course of the proof of this

result, we also show that the upper bound provided by our fluid approximation is asymptotically tight in the same regime. To our knowledge, our approximate policy is the first to yield asymptotic optimality guarantee under dependent demands with arbitrary distributions. The proof for our performance guarantee uses techniques that have not been used in the related literature. In particular, under the standard demand model, to analyze policies from fluid approximations, we upper bound the total consumption of a resource by using a random variable expressed as a sum of independent random variables. In this case, we can use a concentration inequality for sums of independent random variables to upper bound the tail probability of the total consumption of a resource, which in turn, yields a lower bound on the probability that the policy has enough resource availabilities to accept the product request at each time period. The lower bound on the resource availability probabilities are used to lower bound the performance of the policy.

Because the demands in successive stages are dependent in our setting, concentration inequalities for sums of independent random variables are not helpful to us, so we explicitly construct the concentration inequalities that we need. It is standard to use the moment generating function of a random variable to bound its tail probabilities. In particular, if the moment generation function of the random variable Z satisfies $\mathbb{E}\{e^{\lambda Z}\} \leq f(\lambda)$ for all $\lambda \geq 0$, then we can bound its tail probabilities as $\mathbb{P}\{Z \geq c\} = \mathbb{P}\{e^{\lambda Z} \geq e^{\lambda c}\} \leq \frac{1}{e^{\lambda c}} \mathbb{E}\{e^{\lambda Z}\} \leq \frac{f(\lambda)}{e^{\lambda c}}$, where the first inequality is the Markov inequality. We use martingales and the method of bounded differences to bound the moment generating functions of the capacity consumptions of the resources; see Dubhashi and Panconesi (2009). This technique has been used in analyzing randomized algorithms, but their use in revenue management appears to be new. We hope that our use of this technique will further motivate other fluid approximations under even more sophisticated demand models.

Policies under Independent Demands. When the demands in different stages are independent, our revenue management model captures the case where the number of customer arrivals in, say, each week has an arbitrary distribution and the decision maker has a concept of when each week ends and the next one starts. Thus, this demand model is different from a demand model, where there is a single stage consisting of possibly multiple weeks and the number of customer arrivals in the single stage has an arbitrary distribution. In particular, the uncertainty in the demand in our model resolves sequentially over multiple stages. Recalling that the variance proxy of a sub-Gaussian random variable is an upper bound on its variance, under the assumption that the demand in each stage is sub-Gaussian with variance proxy $\sigma^2/200$, we give an approximate policy with a performance guarantee of $\max\left\{\frac{1}{4L}, \left(1 - 4 \frac{\sqrt{(c_{\min} + \sigma^2(K-1)) \log c_{\min}}}{c_{\min}} - \frac{L}{c_{\min}}\right)\right\}$. Variance may be a more intuitive statistic to work with than a lower bound on the mass function.

Even when the numbers of customer arrivals in different stages are independent, establishing the performance guarantee in the previous paragraph requires going one step beyond concentration

inequalities for sums of independent random variables. In particular, the random variable that captures the demand in a particular stage creates dependence between the capacity consumptions of a resource at different time periods in the stage. Thus, as far as we can see, our performance guarantee with independent demands in each stage does not follow from the proof techniques used for the existing asymptotic optimality results under a discrete-time approximation to a Poisson process. We end up constructing the concentration inequalities that we need by exploiting the assumption of sub-Gaussian demand random variables. The assumption of sub-Gaussian demand random variables is relatively mild, as this class is rather general; see Section 2.1.2 in Wainwright (2019). Any bounded random variable, for example, is sub-Gaussian.

Computational Performance. To our knowledge, there is no work on asymptotically tight fluid approximations and asymptotically optimal policies for revenue management problems in which the demands over different time intervals are dependent and have arbitrary distributions. Building such fluid approximations and approximate policies is theoretically interesting, but fluid approximations with a sound theoretical footing can also make a significant impact in practice. In our computational experiments, we make comparisons with existing fluid approximations. While we can show that these fluid approximations do provide upper bounds, they do not provide asymptotically optimal policies. Our fluid approximation, owing to its sound theoretical footing, provides significantly tighter upper bounds and better approximate policies for a range of test problems.

Related Literature: There is significant work on fluid approximations in revenue management, but this work is under demand models that use a discrete-time approximation to a Poisson process, ruling out the possibility of having arbitrary demand distributions and dependence between demands over different time intervals. Considering a revenue management problem with a single resource, Gallego and van Ryzin (1994) show that if we scale the expected demand and the capacity of the resource with the same rate θ , then a policy from a fluid approximation has a performance guarantee of $1 - O(\frac{1}{\sqrt{\theta}})$. Gallego and van Ryzin (1997) generalize this result to a network of resources, where the sale of different products consumes capacities of different combinations of resources. The policies in these papers use the primal solution to a fluid approximation. Talluri and van Ryzin (1998) use the dual solution to construct an asymptotically optimal policy in the same regime. Liu and van Ryzin (2008) and Gallego et al. (2004) construct similar asymptotically optimal policies under customer choice behavior, where the customers choose among the sets of products offered to them. Considering the case where the customers with bookings do not necessarily show up at the time of service, Kunnumkal et al. (2012) give an asymptotically optimal policy that allows overbooking. The papers discussed so far solve the fluid approximation once at the beginning of the selling horizon. Jasin and Kumar (2012) show that solving the fluid approximation periodically over

the selling horizon provides policies with substantially better performance guarantees. Letting c_{\min} be the smallest resource capacity, Rusmevichientong et al. (2020) give a policy with a performance guarantee of $1 - O\left(\frac{1}{\sqrt[3]{c_{\min}}}\right)$. The asymptotic regime in their paper is different in the sense that the authors allow the expected demand to be scaled in an arbitrary fashion. In the same asymptotic regime, Ma et al. (2021), Aouad and Ma (2022), Bai et al. (2022) and Feng et al. (2022) give policies with a performance guarantee of $1 - O\left(\frac{1}{\sqrt{c_{\min}}}\right)$. Aouad and Ma (2022) allow random number of customer arrivals. Balseiro et al. (2023) give a unified analysis for fluid approximations, allowing the possibility of solving the fluid approximation periodically over the selling horizon.

Under a random number of customer arrivals, Bai et al. (2023) give a policy with a performance guarantee of $1 - O\left(\frac{1}{\sqrt{c_{\min}}}\right)$. In their paper, the customer arrivals occur in one stage, so there is no possibility of introducing dependence between the numbers of customer arrivals in successive stages. Thus, the authors do not model the dependence between the numbers of customer arrivals over different time intervals, whereas our focus is to precisely deal with such dependence. Also, even if the demands in successive stages are independent, when the uncertainty in the demand resolves sequentially over multiple stages, we demonstrate that intuitive modifications of the model in Bai et al. (2023) do not yield asymptotically optimal policies or upper bounds. Lastly, the performance guarantee in Bai et al. (2023) directly follows from one-sided Bernstein inequality, whereas we derive our concentration inequalities from scratch. Letting L be the maximum number of resources used by a product, Jiang (2023) gives a policy with a performance guarantee of $\frac{1}{1+L}$ when the distribution governing the products requested by the customers evolves from one time period to the next according to an exogenous Markov chain. The performance guarantee for the policy does not improve when we deal with large systems commanding large product demands and involving large resource capacities. Our performance guarantee converges to one when we deal with large systems. To establish the performance guarantee of $\frac{1}{1+L}$, Jiang (2023) uses a linear program obtained by using linear value function approximations. We show that our fluid approximation is at least as tight as this linear program. Also, Jiang (2023) does not use the linear program to construct a policy, whereas our policy directly uses the optimal solution to our fluid approximation.

Organization: In Section 2, we formulate our revenue management model with arbitrary demand distributions in each stage and dependence between the demands in successive stages. In Section 3, we give the fluid approximation corresponding to our model and show that its optimal objective value is an upper bound on the optimal total expected revenue. In Section 4, we describe the approximate policy from the fluid approximation and give a performance guarantee for the approximate policy. In Section 5, we prove our performance guarantee. In Section 6, we conclude. In the online appendix, we give extended results for our paper, where we focus our results to the case with independent demands in different stages and give computational experiments.

2. Problem Formulation

The set of resources is \mathcal{L} . The capacity of resource i is c_i . The set of products is \mathcal{J} . The revenue of product j is f_j . The resources used by product j are given by the vector $\mathbf{a}_j = (a_{ij} : i \in \mathcal{L}) \in \{0, 1\}^{|\mathcal{L}|}$, where $a_{ij} = 1$ if and only if product j uses resource i . We divide the selling horizon into K stages indexed by $\mathcal{K} = \{1, \dots, K\}$. We use the random variable D^k to capture the number of customer arrivals in stage k . There are at most T customer arrivals in each stage. We divide each stage into T time periods indexed by $\mathcal{T} = \{1, \dots, T\}$. We use λ_{jt}^k to denote the probability that the customer arriving at time period t in stage k requests product j , so we have $\sum_{j \in \mathcal{J}} \lambda_{jt}^k = 1$. We refer D^k as the demand in stage k . The demands in successive stages follow a Markov process. Thus, conditional on D^k , D^{k+1} is independent of D^1, \dots, D^{k-1} . We characterize the evolution of the demands by the survival rate function $\theta_t^k(q) = \mathbb{P}\{D^k \geq t+1 \mid D^k \geq t, D^{k-1} = q\}$, capturing the probability that the demand in stage k is at least $t+1$, given that the demand in stage k is at least t and the demand in the previous stage was q . We assume that $\mathbb{P}\{D^{k+1} = p \mid D^k = q\} \geq \epsilon$ for all $p, q \in \mathcal{T}$ and $k \in \mathcal{K}$ for some $\epsilon > 0$, so the demand in any stage takes values over its full support.

Each stage is a canonical interval of time on the calendar, such as, a day, a week or a month. We are aware of the calendar in the sense that we know when the current stage starts, but we only have probabilistic information about the number of customer arrivals in each stage. Customers in the current stage arrive one by one. Each arriving customer makes a request for a product. We decide whether to accept the product request. Our goal is to find a policy to decide which customer requests to accept so that we maximize the total expected revenue over the selling horizon. We give a dynamic program to compute the optimal policy. We use $\mathbf{y} = (y_i : i \in \mathcal{L}) \in \mathbb{Z}_+^{|\mathcal{L}|}$ to capture the state of the system, where y_i is the remaining capacity of resource i . We use $\mathbf{u} = (u_j : j \in \mathcal{J}) \in \{0, 1\}^{|\mathcal{J}|}$ to capture the decisions, where $u_j = 1$ if and only if we accept a request for product j . The set of feasible decisions is given by $\mathcal{F}(\mathbf{y}) = \{\mathbf{u} \in \{0, 1\}^{|\mathcal{J}|} : a_{ij} u_j \leq y_i \ \forall i \in \mathcal{L}, j \in \mathcal{J}\}$, ensuring that if we want to accept a request for product j and the product uses resource i , then we need to have at least one unit of remaining capacity for resource i . We can find the optimal policy by computing the value functions $\{J_t^k : t \in \mathcal{T}, k \in \mathcal{K}\}$ through the dynamic program

$$J_t^k(\mathbf{y}, q) = \max_{\mathbf{u} \in \mathcal{F}(\mathbf{y})} \left\{ \sum_{j \in \mathcal{J}} \lambda_{jt}^k \left\{ f_j u_j + \theta_t^k(q) J_{t+1}^k(\mathbf{y} - \mathbf{a}_j u_j, q) + (1 - \theta_t^k(q)) J_1^{k+1}(\mathbf{y} - \mathbf{a}_j u_j, t) \right\} \right\}, \quad (1)$$

with the boundary condition that $J_1^{K+1} = 0$. Note that the state variable above keeps both the remaining capacities of the resources and the demand in the previous stage.

In (1), we have a request for product j at time period t in stage k with probability λ_{jt}^k . If we accept this request, then we generate a revenue of f_j and consume the capacities of the resources

used by the product. Given that the demand in the current stage is already t and the demand in the previous stage was q , we have one more demand in the current stage with probability $\theta_t^k(q)$. The demand in the stage right before the beginning of the selling horizon is D^0 and it is fixed at \widehat{D}^0 as problem data. Using $\mathbf{c} = (c_i : i \in \mathcal{L})$ to denote the initial resource capacities, the optimal total expected revenue is $\text{OPT} = J_1^1(\mathbf{c}, \widehat{D}^0)$. The form of the survival rate function is general, so the demand in each stage can have an arbitrary distribution. In our model, we specify the distribution of the demand in the next stage conditional on the demand in the current stage, which implies a certain distribution for the total demand. In Appendix A, we show that we can also calibrate our model to match a given distribution for the total demand. In our model, we decide whether to accept each product request, but our results hold when we make pricing or assortment offering decisions. We will consider such extensions.

3. Fluid Approximation

The dynamic program in (1) involves a high-dimensional state variable, so it is computationally difficult to compute the optimal policy by solving this dynamic program. We give a fluid approximation that will serve two purposes. First, we will use the fluid approximation to obtain an upper bound on the optimal total expected revenue. In this case, we can compare the total expected revenue obtained by any policy with the upper bound to assess the optimality gap of the policy. Second, we will use the fluid approximation to construct a policy that is asymptotically optimal as the capacities of the resources and demands get large. In our fluid approximation, we use the decision variable $x_{jt}^k(q)$ to capture the probability of accepting a request for product j at time period t in stage k given that the demand in the previous stage was q . Using the vector of decision variables $\mathbf{x} = (x_{jt}^k(q) : j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$, to approximate the optimal total expected revenue over the selling horizon, consider linear program

$$\begin{aligned} \bar{Z}_{\text{LP}} = & \max_{\mathbf{x} \in \mathbb{R}_+^{|\mathcal{J}|T^2K}} \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{q \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j \mathbb{P}\{D^k \geq t, D^{k-1} = q\} x_{jt}^k(q) \\ \text{st} & \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\} x_{js}^\ell(p) \\ & + \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} x_{js}^k(q) \leq c_i \quad \forall i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K} \\ & x_{jt}^k(q) \leq \lambda_{jt}^k \quad \forall j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K}. \end{aligned} \tag{2}$$

In the linear program above, the objective function accounts for the total expected revenue over the selling horizon. In particular, we can make a sale for product j at time period t in stage k

only if the demand in stage k is at least t . Furthermore, if the demand in stage $k - 1$ is q , then we make a sale for product j at time period t in stage k with probability $x_{jt}^k(q)$. Therefore, the expression $\sum_{q \in \mathcal{T}} \mathbb{P}\{D^k \geq t, D^{k-1} = q\} x_{jt}^k(q)$ corresponds to the expected sales for product j at time period t in stage k . The left side of the first constraint corresponds to the total expected capacity consumption of resource i up to and including time period t in stage k *conditional* on the fact that the demand in stage k is at least t and the demand in stage $k - 1$ is q . In the first sum, similar to the objective function, the expression $\sum_{p \in \mathcal{T}} \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\} x_{js}^\ell(p)$ corresponds to the expected sales for product j at time period s in stage ℓ conditional on the fact that the demand in stage k is at least t and the demand in stage $k - 1$ is q . If product j uses resource i , then these sales consume the capacity of resource i . In the second sum, conditional on the fact that the demand in stage k is at least t and the demand in stage $k - 1$ is q , we can make a sale for product j at all time periods in stage k up to and including time period t . Furthermore, we accept a request for product j at time period s in stage k with probability $x_{js}^k(q)$. Thus, the first constraint is the capacity constraint. The second constraint is the demand constraint, ensuring that the probability of accepting a request for a product at any time period in any stage does not exceed the probability of getting the request. We emphasize two novel aspects of the linear program given above.

First, because the demand in stage k depends on the demand in stage $k - 1$, the probability of accepting a request for a product at any time period in stage k depends on the demand in stage $k - 1$ as well. Second, perhaps more surprisingly, the first constraint keeps track of the total expected capacity consumption of a resource up to and including time period t in stage k , conditional on the fact that the demand in stage k is at least t and the demand in stage $k - 1$ is q . The form of this conditioning is unexpected. We can use the Markovian structure of the demands to slightly simplify the first constraint. If $\ell \leq k - 1$, then conditional on D^{k-1} , D^ℓ and $D^{\ell-1}$ are independent of D^k . Thus, we can replace the probability $\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\}$ in the first sum with $\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^{k-1} = q\}$. Furthermore, the second sum is increasing in t , so we can replace the sum $\sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} x_{js}^k(q)$ with $\sum_{s \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} x_{js}^k(q)$. Therefore, we can express the first constraint equivalently as $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^{k-1} = q\} x_{js}^\ell(p) + \sum_{s \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} x_{js}^k(q) \leq c_i$ for all $i \in \mathcal{L}$, $q \in \mathcal{T}$ and $k \in \mathcal{K}$. In this way, we can reduce the number of constraints in the first constraint by a factor of T . Nevertheless, we believe that our fluid approximation, as stated in (2), is more instructive, so we keep it in its full form. Lastly, we can write $\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^{k-1} = q\} = \mathbb{P}\{D^{\ell-1} = p \mid D^\ell \geq s, D^{k-1} = q\} \mathbb{P}\{D^\ell \geq s \mid D^{k-1} = q\}$, but it is not true that $\mathbb{P}\{D^{\ell-1} = p \mid D^\ell \geq s, D^{k-1} = q\} = \mathbb{P}\{D^{\ell-1} = p \mid D^\ell \geq s\}$, so there is no more simplification.

Our demand model is a generalization of the standard demand model based on a discrete-time approximation to a Poisson process. If there is one stage in the selling horizon and the number

of time periods in the stage takes a deterministic value, then our demand model reduces to the standard demand model based on a discrete-time approximation to a Poisson process. Accordingly, if there is one stage in the selling horizon and the number of time periods in the stage takes a deterministic value, then problem (2) reduces to the standard fluid approximation. In particular, if there is one stage and the number of time periods in the stage takes the deterministic value of T , then we have $\mathbb{P}\{D^1 \geq t, D^0 = q\} = 1$ for all $t = 1, \dots, T$ and $q = \widehat{D}^0$, whereas $\mathbb{P}\{D^1 \geq t, D^0 = q\} = 0$ for all $t = 1, \dots, T$ and $q \neq \widehat{D}^0$. Thus, noting that there is one stage, only the decision variables $(x_{jt}^1(\widehat{D}^0) : j \in \mathcal{J}, t \in \mathcal{T})$ have a non-zero objective function coefficient in (2), so we can drop all other decision variables. In this case, using y_{jt} to succinctly denote the decision variable $x_{jt}^1(\widehat{D}^0)$, the objective function in (2) becomes $\sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j y_{jt}$. Because there is one stage, the first sum in the first constraint with $k = 1$ is zero, so noting that we drop the decision variables other than $(x_{jt}^1(\widehat{D}^0) : j \in \mathcal{J}, t \in \mathcal{T})$, the first constraint in (2) becomes $\sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} y_{js} \leq c_i$ for all $i \in \mathcal{L}$ and $t \in \mathcal{T}$. Lastly, the second constraint in (2) becomes $y_{jt} \leq \lambda_{jt}^1$ for all $j \in \mathcal{J}$ and $t \in \mathcal{T}$. Therefore, the linear program in (2) reduces to the standard fluid approximation.

Upper Bound on the Optimal Total Expected Revenue:

Turning back to our demand model with multiple stages in the selling horizon with the demand in each stage having an arbitrary distribution, possibly with dependence between the demands in different stages, we show that the optimal objective value of the linear program in (2) is an upper bound on the optimal total expected revenue. There are two uses of this result. From the practical side, it is difficult to compute the optimal policy, but we can compare the performance of any heuristic policy with the upper bound on the optimal total expected revenue to assess the optimality gap of the heuristic policy. If the total expected revenue of the heuristic policy is close to the upper bound, then we can conclude that the heuristic policy is satisfactory. From the theoretical side, we will give performance guarantees for a policy that is obtained by using the linear program in (2). Because it is difficult to compute the optimal policy, we will establish these performance guarantees by comparing the total expected revenue of the policy with the upper bound on the optimal total expected revenue. Thus, the linear program in (2) also becomes useful to establish performance guarantees. In the next theorem, we show that the optimal objective value of problem (2) is indeed an upper bound on the optimal total expected revenue.

Theorem 3.1 (Upper Bound) *Using OPT to denote the optimal total expected revenue and \overline{Z}_{LP} to denote the optimal objective value of problem (2), we have $\overline{Z}_{\text{LP}} \geq \text{OPT}$.*

The proof of the theorem, which is based on relaxing the resource availability constraints in (1) through Lagrange multipliers, is in Appendix B. There are other possible approaches to show

that problem (2) provides an upper bound on the optimal total expected revenue. In Appendix C, we use linear approximations to the value functions. To compute the slopes and intercepts of the approximations, we plug the approximations into the linear programming formulation of the dynamic program in (1). A suitable relaxation of this linear program is the dual of problem (2). Jiang (2023) uses a similar outline to develop a linear program that provides an upper bound. In the development in the appendix, we also show that our relaxation is at least as tight as the one in Jiang (2023). In Appendix D, we use the decisions made by the optimal policy to construct a feasible solution to problem (2). This approach is not constructive in the sense that it does not allow us to derive the form of problem (2), but it simply verifies that the optimal objective value of problem (2) is an upper bound. Lastly, another approach for obtaining an upper bound on the optimal total expected revenue is to use an offline bound, where we observe the realizations of all product requests and solve an optimization problem to choose the product requests to accept. Under the standard demand model with no demand dependence, the offline bound is at least as tight as the one provided by the fluid approximation; see Talluri and van Ryzin (1999). The analogue of this result does not hold under our demand model with demand dependence. In Appendix E, we give problem instances to demonstrate that the offline bound can be tighter or looser than the optimal objective value of problem (2).

4. Approximate Policy and Asymptotic Optimality

We use an optimal solution to problem (2) to construct an approximate policy. We show that this policy has a constant-factor performance guarantee, but if both the number of stages in the selling horizon and capacities of the resources increase with the same rate, then the policy is asymptotically optimal. Thus, we expect the approximate policy to perform particularly well under large product demands and large resource capacities, but because of its constant-factor guarantee, the approximate policy can never perform arbitrarily badly. In our approximate policy, we solve the linear program in (2) once at the beginning of the selling horizon. Letting $\bar{x} = (\bar{x}_{jt}^k(q) : j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$ be an optimal solution, we make the decisions as follows.

Approximate Policy from the Fluid Approximation:

Using $\gamma \in [0, 1]$ to denote a tuning parameter, if we have a request for product j at time period t in stage k and the demand in stage $k - 1$ was q , then we are willing to accept the request with probability $\gamma \frac{\bar{x}_{jt}^k(q)}{\lambda_{jt}^k}$. If we are willing to accept the request and there are enough resource capacities to accept the request, then we accept the request. Otherwise, we reject.

The description of the approximate policy as given above will be adequate to establish its performance guarantee. In Appendix F, we give a detailed description of our approximate policy.

Letting $c_{\min} = \min_{i \in \mathcal{L}} c_i$ be the smallest resource capacity and $L = \max_{j \in \mathcal{J}} \sum_{i \in \mathcal{L}} a_{ij}$ be the maximum number of resources used by a product, we have the next performance guarantee.

Theorem 4.1 (Performance Guarantee) *Using APX to denote the total expected revenue of the approximate policy, there exists a choice of the tuning parameter γ such that we have*

$$\frac{\text{APX}}{\text{OPT}} \geq \frac{\text{APX}}{\overline{Z}_{\text{LP}}} \geq \max \left\{ \frac{1}{4L}, \left(1 - 4 \frac{\sqrt{(c_{\min} + \frac{1}{\epsilon^6} (K-1)) \log c_{\min}}}{c_{\min}} - \frac{L}{c_{\min}} \right) \right\}.$$

We devote the next section to the proof of the theorem. To our knowledge, this theorem gives the first policy with an asymptotic performance guarantee under dependent demands. The proof involves novel ideas. We use the moment generating function of the resource capacity consumptions to lower bound the probability that we have enough resource capacities to accept a product request. Because of the dependence between the demands in different stages, it is difficult to characterize the moment generation function of the resource capacity consumptions. We bound the moment generating functions by using martingales and the method of bounded differences; see Chapter 5 in Dubhashi and Panconesi (2009). Dependence between the demands requires us to derive our own moment generating function bounds, which ultimately yield the tail probability bounds needed for our performance guarantee. If there is one stage in the selling horizon and the number of time periods in the stage takes a deterministic value, then our demand model reduces to the standard demand model based on a discrete-time approximation to a Poisson process. Other papers, such as Ma et al. (2020) and Baek and Ma (2022), focus on performance guarantees under the standard demand model. Our demand model is a generalization of the standard demand model, so Theorem 4.1 holds under the standard demand model. Setting $K = 1$ in Theorem 4.1, our approximate policy has a performance guarantee of $\max \left\{ \frac{1}{4L}, 1 - 4 \sqrt{\frac{\log c_{\min}}{c_{\min}}} - \frac{L}{c_{\min}} \right\}$ under the standard demand model. Feng et al. (2022) also study fluid approximations under the standard demand model.

We proceed to interpreting the two parts in the performance guarantee in Theorem 4.1. In many network revenue management settings, the number of resources and number of products can be large, but the number of resources used by a particular product remains bounded. In the airline setting, for example, we may have hundreds of flight legs and thousands of itineraries, but the number of flight legs in an itinerary rarely exceeds two, corresponding to $L = 2$. Thus, the first part in the performance guarantee provides a constant-factor performance guarantee for the approximate policy when L is uniformly bounded. On the other hand, consider a regime where we scale both the number of stages and resource capacities with the same rate θ , so that $K = \theta \overline{K}$ and $c_{\min} = \theta \overline{c}$ for some fixed $\overline{K}, \overline{c} \in \mathbb{Z}_+$. If θ gets large, then the expected demands for the products and capacities for the resources both get large. Letting APX^θ be the total expected revenue from the

approximate policy, OPT^θ be the optimal total expected revenue and $\bar{Z}_{\text{LP}}^\theta$ be the optimal objective value of problem (2) when we scale the number of stages and resource capacities with θ , by Theorem 4.1, we have $1 \geq \frac{\text{APX}^\theta}{\text{OPT}^\theta} \geq \frac{\text{APX}^\theta}{\bar{Z}_{\text{LP}}^\theta} \geq 1 - \frac{4}{\epsilon^3} \frac{\sqrt{(\bar{c}+K) \log(\theta \bar{c})}}{\bar{c}\sqrt{\theta}} - \frac{L}{\theta \bar{c}}$. In this case, ignoring the logarithmic terms, as θ gets large, the relative gap between the total expected revenue of the approximate policy and the optimal total expected revenue converges to one with rate $1 - \frac{1}{\sqrt{\theta}}$. Therefore, as the number of stages in the selling horizon and capacities of the resources increase with the same rate, the approximate policy is asymptotically optimal. Similarly, as θ gets large, the relative gap between the total expected revenue of the approximate policy and the optimal objective value of the fluid approximation converges to one with rate $1 - \frac{1}{\sqrt{\theta}}$ as well. In Appendix G, we also give a problem instance such that if we scale the number of stages and resource capacities with the same rate θ , then the optimal total expected revenue and the optimal objective value of problem (2) satisfy $\frac{\text{OPT}^\theta}{\bar{Z}_{\text{LP}}^\theta} \leq 1 - \frac{1}{\sqrt{26\theta}}$. In other words, the optimal total expected revenue is at most within a factor of $O\left(1 - \frac{1}{\sqrt{\theta}}\right)$ of the optimal objective value of the fluid approximation. Therefore, we cannot improve the asymptotic performance guarantee for the approximate policy as long as we establish this performance guarantee by comparing the total expected revenue of the approximate policy with the optimal objective value of the fluid approximation. In this sense, the performance guarantee that we give for the approximate policy is asymptotically tight.

The scaling regime in the previous paragraph increases the number of stages and resource capacities with the same rate. Because the demands in different stages are dependent, increasing the number of stages in the selling horizon is perhaps the most natural approach to increase the expected demands for the products. In this way, we can increase the expected demands for the products without distorting the correlation structure for the demands in different stages. Another approach to increase the expected demands for the products could be to increase the support of the demand in each stage, while keeping the number of stages in the selling horizon fixed. Scaling the expected demand in this fashion can potentially distort the correlation structure for the demands in different stages. In Appendix H, we also give a counterexample to demonstrate that the relative gap between the total expected revenue of the approximate policy and the optimal objective value of the fluid approximation in (2) does not necessarily converge to one as we increase the support of the demand in each stage and capacities of the resources with the same rate, while keeping the number of stages constant. In our counterexample, we give a problem instance with three stages. There is a single resource with a capacity of $C + 1$. The largest value of the demand in a week is C . There are two products. We show that the optimal total expected revenue is C , whereas the optimal objective value of problem (2) is $\frac{5}{4}C$. Thus, we have $\frac{\text{APX}^\theta}{\bar{Z}_{\text{LP}}^\theta} \leq \frac{\text{OPT}^\theta}{\bar{Z}_{\text{LP}}^\theta} = \frac{4}{5}$. For this problem instance, no matter how large C is, the ratio $\frac{\text{APX}^\theta}{\bar{Z}_{\text{LP}}^\theta}$ always stays away from one.

5. Performance Guarantee

In this section, we give a proof for Theorem 4.1. We focus on showing the performance guarantee $\frac{\text{APX}}{\text{OPT}} \geq \frac{\text{APX}}{\bar{Z}_{\text{LP}}} \geq 1 - 4 \frac{\sqrt{(c_{\min} + \frac{1}{e^6}(K-1)) \log c_{\min}}}{c_{\min}} - \frac{L}{c_{\min}}$. To establish this performance guarantee, we use ideas that have not been used in the revenue management literature to analyze fluid approximations. In Appendix I, we turn our attention to showing the performance guarantee $\frac{\text{APX}}{\text{OPT}} \geq \frac{\text{APX}}{\bar{Z}_{\text{LP}}} \geq \frac{1}{4L}$, which is more straightforward. A common approach for analyzing approximate policies from fluid approximations involves using an auxiliary random variable to upper bound the capacity consumption of a resource under the approximate policy. Thus, we can use a concentration inequality to upper bound the tail probabilities of the auxiliary random variable, in which case, we can lower bound the probability that there is enough capacity to accept different product requests at different time periods in the selling horizon; see, for example, Feng et al. (2022). This approach usually exploits the fact that the auxiliary random variable can be expressed as a sum of independent random variables, which facilitates using concentration inequalities for sums of independent random variables. Because the demands in different stages are dependent in our setting, we cannot construct similar auxiliary random variables that can be expressed as sums of independent random variables. Thus, we resort to new ideas.

Preliminary Random Variables and Availability Probabilities:

We define four classes of Bernoulli random variables for each $k \in \mathcal{K}$ and $t \in \mathcal{T}$. Analogues of these random variables have been used in the analysis of other fluid approximations.

- *Demand in Each Stage.* For each $q \in \mathcal{T}$, the random variable $\Psi_t^k(q)$ takes value one if we reach time period t in stage k before this stage is over and the demand in stage $k-1$ is q . In other words, letting $\mathbf{1}(\cdot)$ be the indicator function, $\Psi_t^k(q) = \mathbf{1}(D^k \geq t, D^{k-1} = q)$.

- *Product Request.* For each $j \in \mathcal{J}$, the random variable A_{jt}^k takes value one if the customer arriving at time period t in stage k requests product j . We have $\mathbb{P}\{A_{jt}^k = 1\} = \lambda_{jt}^k$. The random variables $\{A_{jt}^k : t \in \mathcal{T}, k \in \mathcal{K}\}$ are independent of each other.

- *Policy Decision.* For each $j \in \mathcal{J}$ and $q \in \mathcal{T}$, the random variable $X_{jt}^k(q)$ takes value one if the approximate policy is willing to accept a request for product j at time period t in stage k when the demand in stage $k-1$ was q . By our approximate policy, $\mathbb{P}\{X_{jt}^k(q) = 1\} = \gamma \frac{\bar{x}_{jt}^k(q)}{\lambda_{jt}^k}$.

- *Availability.* For each $j \in \mathcal{J}$, the random variable G_{jt}^k takes value one if we have enough capacity to accept a request for product j at time period t in stage k under the approximate policy. Instead of calculating the probability $\mathbb{P}\{G_{jt}^k = 1\}$, we will lower bound $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\}$.

The random variables A_{jt}^k and $X_{jt}^k(q)$ are both simple Bernoulli draws independent of the decisions of the approximate policy, remaining capacities of the resources or realizations of the

demands in different stages. Under the approximate policy, the sales for product j at time period t in stage k is given by $\sum_{q \in \mathcal{T}} \Psi_t^k(q) G_{jt}^k A_{jt}^k X_{jt}^k(q)$, where we use the fact that we sell product j at time period t in stage k if we reach time period t in stage k , there is enough capacity to accept a request for product j , we have a request for the product and the approximate policy is willing to accept the product request. The remaining capacities of the resources at time period t in stage k depend on the requests for the products and decisions of the approximate policy at the earlier time periods, but not at time period t in stage k . Thus, taking expectations, the expected sales for product j at time period t in stage k is $\sum_{q \in \mathcal{T}} \mathbb{P}\{\Psi_t^k(q) = 1\} \mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} \mathbb{P}\{A_{jt}^k = 1\} \mathbb{P}\{X_{jt}^k(q) = 1\}$. In this case, noting that $\mathbb{P}\{A_{jt}^k = 1\} = \lambda_{jt}^k$ and $\mathbb{P}\{X_{jt}^k(q) = 1\} = \gamma \frac{\bar{x}_{jt}^k(q)}{\lambda_{jt}^k}$, we can write the last expectation equivalently as $\sum_{q \in \mathcal{T}} \mathbb{P}\{D^k \geq t, D^{k-1} = q\} \mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} \gamma \bar{x}_{jt}^k(q)$, capturing the expected sales for product j at time period t in stage k under the approximate policy. Thus, the total expected revenue of the approximate policy is given by

$$\text{APX} = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \sum_{q \in \mathcal{T}} f_j \mathbb{P}\{D^k \geq t, D^{k-1} = q\} \mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} \gamma \bar{x}_{jt}^k(q). \quad (3)$$

By the definition of $\bar{\mathbf{x}}$, we have $\bar{Z}_{\text{LP}} = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \sum_{q \in \mathcal{T}} f_j \mathbb{P}\{D^k \geq t, D^{k-1} = q\} \bar{x}_{jt}^k(q)$, so if we can show that $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} \geq \alpha$, then we get $\text{APX} \geq \gamma \alpha \bar{Z}_{\text{LP}}$.

Motivated by the discussion in the previous paragraph, we focus on lower bounding the availability probability $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\}$. Under the approximate policy, the sales for product j at time period t in stage k is $\sum_{q \in \mathcal{T}} \Psi_t^k(q) G_{jt}^k A_{jt}^k X_{jt}^k(q)$, so $\sum_{q \in \mathcal{T}} \Psi_t^k(q) A_{jt}^k X_{jt}^k(q)$ is an upper bound on these sales for product j at time period t in stage k . In this case, $\sum_{j \in \mathcal{J}} \sum_{q \in \mathcal{T}} a_{ij} \Psi_t^k(q) A_{jt}^k X_{jt}^k(q)$ is an upper bound on the capacity consumption of resource i at time period t in stage k . Letting $N_{it}^k(q) = \sum_{j \in \mathcal{J}} a_{ij} A_{jt}^k X_{jt}^k(q)$, we express our upper bound on the capacity consumption of resource i at time period t in stage k as $\sum_{q \in \mathcal{T}} \Psi_t^k(q) N_{it}^k(q)$. Note that $\{N_{it}^k(q) : t \in \mathcal{T}, k \in \mathcal{K}\}$ are Bernoulli random variables and they are independent of each other. Having $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{q \in \mathcal{T}} \Psi_s^\ell(q) N_{is}^\ell(q) + \sum_{s=1}^t \sum_{q \in \mathcal{T}} \Psi_s^k(q) N_{is}^k(q) < c_i$ implies that the total capacity consumption of resource i up to and including time period t in stage k does not exceed the capacity of the resource, in which case, we have capacity available for resource i at time period t in stage k . Therefore, letting $\mathcal{L}_j = \{i \in \mathcal{L} : a_{ij} = 1\}$ to capture the set of resources used by product j , we obtain

$$\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} \geq \mathbb{P}\left\{ \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p) + \sum_{s=1}^t \sum_{p \in \mathcal{T}} \Psi_s^k(p) N_{is}^k(p) < c_i \quad \forall i \in \mathcal{L}_j \mid \Psi_t^k(q) = 1 \right\},$$

where we use the fact if the upper bounds on the consumption of the resources used by product j do not exceed their capacities, then we have capacities to accept a request for product j .

By the definition of $\Psi_t^k(q)$, having $\Psi_t^k(q) = 1$ is equivalent to having $D^k \geq t$ and $D^{k-1} = q$. Thus, if $\Psi_t^k(q) = 1$, then we have $D^k \geq s$ for all $s = 1, \dots, t$, $D^{k-1} = q$ and $D^{k-1} \neq p$ for all $p \in \mathcal{T} \setminus \{q\}$.

Therefore, if $\Psi_t^k(q) = 1$, then $\Psi_s^k(q) = 1$ for all $s = 1, \dots, t$ and $\Psi_s^k(p) = 0$ for all $p \in \mathcal{T} \setminus \{q\}$. In this case, the inequality above is equivalent to

$$\begin{aligned}
\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} &\geq \mathbb{P}\left\{\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p) + \sum_{s=1}^t N_{is}^k(q) < c_i \ \forall i \in \mathcal{L}_j \mid \Psi_t^k(q) = 1\right\} \\
&= 1 - \mathbb{P}\left\{\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p) + \sum_{s=1}^t N_{is}^k(q) \geq c_i \text{ for some } i \in \mathcal{L}_j \mid \Psi_t^k(q) = 1\right\} \\
&\geq 1 - \sum_{i \in \mathcal{L}_j} \mathbb{P}\left\{\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p) + \sum_{s=1}^t N_{is}^k(q) \geq c_i \mid \Psi_t^k(q) = 1\right\} \\
&= 1 - \sum_{i \in \mathcal{L}_j} \mathbb{P}\left\{\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p) + \sum_{s=1}^t N_{is}^k(q) \geq c_i \mid D^{k-1} = q\right\}, \tag{4}
\end{aligned}$$

where the second inequality is the union bound and the second equality holds because given D^{k-1} , D^1, \dots, D^{k-1} are independent of D^k . Thus, it is enough to upper bound the last probability in (4).

Moment Generating Function Bounds:

The discussion so far in this section has been following standard arguments, but we proceed to introducing new ideas. To upper bound the last probability on the right side of (4), letting $n_{it}^k(q) = \mathbb{E}\{N_{it}^k(q)\}$, for all $i \in \mathcal{L}$ and $k \in \mathcal{K}$, we define $U_i^k = \sum_{\ell=1}^k \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p)$ and $V_i^k = \sum_{\ell=1}^k \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) n_{is}^\ell(p)$. Using the vector of random variables $\mathbf{D}^{[\ell, k]} = (D^\ell, \dots, D^k)$ for notational brevity, note that the random variable V_i^k is a deterministic function of $\mathbf{D}^{[1, k]}$. Because e^x is convex in x , using the Jensen inequality, it is simple to show that $\mathbb{E}\{e^{\lambda V_i^k} \mid D^k\} \leq \mathbb{E}\{e^{\lambda U_i^k} \mid D^k\}$ for all $\lambda \geq 0$, so the moment generating function of the random variable U_i^k conditional on D^k upper bounds its counterpart for the random variable V_i^k . In the next lemma, we characterize the gap between the two moment generating functions.

Lemma 5.1 (Moment Generating Function Gap) *For all $k \in \mathcal{K}$, $i \in \mathcal{L}$ and $\lambda \geq 0$, we have $\mathbb{E}\{e^{\lambda U_i^k} \mid D^k\} \leq e^{\frac{1}{2\epsilon} k \lambda^2} \mathbb{E}\{e^{\lambda V_i^k} \mid D^k\}$.*

Proof: The random variables $\{N_{is}^\ell(p) : s \in \mathcal{T}, \ell = 1, \dots, k\}$ are independent of each other. Also, for $\ell = 1, \dots, k$ and $s, p \in \mathcal{T}$, $\Psi_s^\ell(p)$ is a deterministic function of $\mathbf{D}^{[1, k]}$. Thus, we have

$$\begin{aligned}
\mathbb{E}\{e^{\lambda(U_i^k - V_i^k)} \mid \mathbf{D}^{[1, k]}\} &= \mathbb{E}\{e^{\lambda \sum_{\ell=1}^k \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) (N_{is}^\ell(p) - n_{is}^\ell(p))} \mid \mathbf{D}^{[1, k]}\} \\
&= \prod_{\ell=1}^k \prod_{s \in \mathcal{T}} \mathbb{E}\{e^{\lambda \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) (N_{is}^\ell(p) - n_{is}^\ell(p))} \mid \mathbf{D}^{[1, k]}\}. \tag{5}
\end{aligned}$$

Because $\sum_{p \in \mathcal{T}} \Psi_s^\ell(p) \leq 1$ and the random variable $N_{is}^\ell(p)$ is Bernoulli with expectation $n_{is}^\ell(p)$, we have $\sum_{p \in \mathcal{T}} \Psi_s^\ell(p) (N_{is}^\ell(p) - n_{is}^\ell(p)) \in [-1, 1]$. If the mean-zero random variable Z is bounded by

$[a, b]$, then we have $\mathbb{E}\{e^{\lambda Z}\} \leq e^{\frac{1}{8}(b-a)^2 \lambda^2}$ for any $\lambda \geq 0$; see Lemma 5.1 in Dubhashi and Panconesi (2009). Thus, the last conditional expectation on the right side of (5) is upper bounded by $e^{\frac{1}{2}\lambda^2}$, so by (5), we obtain $\mathbb{E}\{e^{\lambda(U_i^k - V_i^k)} \mid \mathbf{D}^{[1,k]}\} \leq e^{\frac{1}{2}kT\lambda^2}$. The random variable V_i^k is a deterministic function of $\mathbf{D}^{[1,k]}$. In this case, using the tower property of conditional expectations, we get $\mathbb{E}\{e^{\lambda U_i^k} \mid D^k\} = \mathbb{E}\{\mathbb{E}\{e^{\lambda V_i^k} e^{\lambda(U_i^k - V_i^k)} \mid \mathbf{D}^{[1,k]}\} \mid D^k\} = \mathbb{E}\{e^{\lambda V_i^k} \mathbb{E}\{e^{\lambda(U_i^k - V_i^k)} \mid \mathbf{D}^{[1,k]}\} \mid D^k\}$. Using the fact that $\mathbb{E}\{e^{\lambda(U_i^k - V_i^k)} \mid \mathbf{D}^{[1,k]}\} \leq e^{\frac{1}{2}kT\lambda^2}$, we obtain $\mathbb{E}\{e^{\lambda U_i^k} \mid D^k\} \leq e^{\frac{1}{2}kT\lambda^2} \mathbb{E}\{e^{\lambda V_i^k} \mid D^k\}$ by the last chain of equalities. The result follows by noting the assumption that $\mathbb{P}\{D^{k+1} = p \mid D^k = q\} \geq \epsilon$ for all $p, q \in \mathcal{T}$, so $1 = \sum_{p \in \mathcal{T}} \mathbb{P}\{D^{k+1} = p \mid D^k = q\} \geq T\epsilon$, which implies that $T \leq 1/\epsilon$. \blacksquare

If the random variable Z satisfies $\mathbb{E}\{e^{\lambda Z}\} \leq f(\lambda)$ for all $\lambda \geq 0$, then we can upper bound its tail probabilities as $\mathbb{P}\{Z \geq c\} = \mathbb{P}\{e^{\lambda Z} \geq e^{\lambda c}\} \leq \frac{1}{e^{\lambda c}} \mathbb{E}\{e^{\lambda Z}\} \leq \frac{f(\lambda)}{e^{\lambda c}}$, where the first inequality uses the Markov inequality. In the last probability in (4), we have $U_i^{k-1} = \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p)$, so we may use the moment generation function of U_i^k to lower bound the availability probabilities. By Lemma 5.1, the moment generating function of V_i^k can be a proxy for the moment generating function of U_i^k . For all $i \in \mathcal{L}$, $k \in \mathcal{K}$ and $\ell = 1, \dots, k$, we define $M_i^k(\ell) = \mathbb{E}\{V_i^k \mid \mathbf{D}^{[\ell,k]}\}$. Therefore, the random variable $M_i^k(\ell)$ is a deterministic function of $\mathbf{D}^{[\ell,k]}$. Noting that V_i^k is a deterministic function of $\mathbf{D}^{[1,k]}$, we have $M_i^k(1) = \mathbb{E}\{V_i^k \mid \mathbf{D}^{[1,k]}\} = V_i^k$ with probability one.

In the next lemma, we upper bound the moment generating function of $M_i^k(\ell)$ for all $\ell = 1, \dots, k$, which, noting that $M_i^k(1) = V_i^k$, will yield an upper bound on the same for V_i^k .

Lemma 5.2 (Moment Generating Function Bound) *Letting $M_i^0(0) = 0$, for all $k \in \mathcal{K}$, $i \in \mathcal{L}$, $\ell = 0, \dots, k-1$ and $\lambda \geq 0$, we have*

$$\mathbb{E}\{e^{\lambda(M_i^{k-1}(\ell) + \sum_{s=1}^t n_{is}^k(q))} \mid D^{k-1} = q\} \leq e^{\frac{2}{\epsilon^6} (k-1-\ell)\lambda^2 + \gamma c_i \lambda}.$$

Proof: We show the result by using induction over $\ell = 1, \dots, k-1$. Consider the case $\ell = k-1$. We have $\mathbb{E}\{M_i^{k-1}(k-1) \mid D^{k-1}\} = \mathbb{E}\{\mathbb{E}\{V_i^{k-1} \mid D^{k-1}\} \mid D^{k-1}\} = \mathbb{E}\{V_i^{k-1} \mid D^{k-1}\}$, so we get

$$\begin{aligned} \mathbb{E}\{M_i^{k-1}(k-1) \mid D^{k-1} = q\} + \sum_{s=1}^t n_{is}^k(q) &= \mathbb{E}\{V_i^{k-1} \mid D^{k-1} = q\} + \sum_{s=1}^t n_{is}^k(q) \\ &\stackrel{(a)}{=} \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^{k-1} = q\} n_{is}^\ell(p) + \sum_{s=1}^t n_{is}^k(q) \\ &\stackrel{(b)}{=} \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^{k-1} = q\} \gamma \bar{x}_{js}^\ell(p) + \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} \gamma \bar{x}_{js}^k(q) \stackrel{(c)}{\leq} \gamma c_i, \end{aligned}$$

where (a) uses the definition of V_i^{k-1} , (b) follows because we have $n_{is}^\ell(p) = \mathbb{E}\{N_{is}^\ell(p)\}$, in which case, by the definition of $N_{is}^\ell(p)$, we get $n_{is}^\ell(p) = \mathbb{E}\{N_{is}^\ell(p)\} = \sum_{j \in \mathcal{J}} a_{ij} \gamma \bar{x}_{js}^\ell(p)$ and (c) holds by

the first constraint in problem (2), as well as noting that conditional on D^{k-1} , D^ℓ is independent of D^k for $\ell \leq k-1$, so $\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\} = \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^{k-1} = q\}$. By its definition, $M_i^{k-1}(k-1)$ is a deterministic function of D^{k-1} , which implies that given $D^{k-1} = q$, $M_i^{k-1}(k-1)$ is a deterministic quantity. Thus, using the chain of inequalities above, we obtain $\mathbb{E}\{e^{\lambda(M_i^{k-1}(k-1) + \sum_{s=1}^t n_{is}^k(q))} \mid D^{k-1} = q\} = e^{\mathbb{E}\{\lambda(M_i^{k-1}(k-1) + \sum_{s=1}^t n_{is}^k(q)) \mid D^{k-1} = q\}} \leq e^{\gamma c_i \lambda}$, so the result holds for $\ell = k-1$. Assuming that the result holds for $\ell+1 \leq k-1$, we show that the result holds for $\ell \leq k-1$. By Lemma J.3 in Appendix J, $|M_i^{k-1}(\ell) - M_i^{k-1}(\ell+1)| \leq \frac{2}{\epsilon^3}$ with probability one. Also, by the tower property of conditional expectations, using the definition of $M_i^{k-1}(\ell)$, we have $\mathbb{E}\{M_i^{k-1}(\ell) \mid \mathbf{D}^{[\ell+1, k-1]}\} = \mathbb{E}\{\mathbb{E}\{V_i^{k-1} \mid \mathbf{D}^{[\ell, k-1]}\} \mid \mathbf{D}^{[\ell+1, k-1]}\} = \mathbb{E}\{V_i^{k-1} \mid \mathbf{D}^{[\ell+1, k-1]}\}$. Using precisely the same argument, we can verify that $\mathbb{E}\{M_i^{k-1}(\ell+1) \mid \mathbf{D}^{[\ell+1, k-1]}\} = \mathbb{E}\{V_i^{k-1} \mid \mathbf{D}^{[\ell+1, k-1]}\}$ as well. Thus, conditional on $\mathbf{D}^{[\ell+1, k-1]}$, the random variable $M_i^{k-1}(\ell) - M_i^{k-1}(\ell+1)$ is mean-zero and bounded by $[-\frac{2}{\epsilon^3}, \frac{2}{\epsilon^3}]$. Recall that if the mean-zero random variable Z is bounded by $[a, b]$, then we have $\mathbb{E}\{e^{\lambda Z}\} \leq e^{\frac{1}{8}(b-a)^2 \lambda^2}$. In this case, we get $\mathbb{E}\{e^{\lambda(M_i^{k-1}(\ell) - M_i^{k-1}(\ell+1))} \mid \mathbf{D}^{[\ell+1, k-1]}\} \leq e^{\frac{2}{\epsilon^6} \lambda^2}$. Thus, using the fact that $M_i^{k-1}(\ell+1)$ is a deterministic function of $\mathbf{D}^{[\ell+1, k-1]}$, we have

$$\begin{aligned} \mathbb{E}\{e^{\lambda M_i^{k-1}(\ell)} \mid D^{k-1}\} &= \mathbb{E}\{\mathbb{E}\{e^{\lambda M_i^{k-1}(\ell+1)} e^{\lambda(M_i^{k-1}(\ell) - M_i^{k-1}(\ell+1))} \mid \mathbf{D}^{[\ell+1, k-1]}\} \mid D^{k-1}\} \\ &= \mathbb{E}\{e^{\lambda M_i^{k-1}(\ell+1)} \mathbb{E}\{e^{\lambda(M_i^{k-1}(\ell) - M_i^{k-1}(\ell+1))} \mid \mathbf{D}^{[\ell+1, k-1]}\} \mid D^{k-1}\} \leq e^{\frac{2}{\epsilon^6} \lambda^2} \mathbb{E}\{e^{\lambda M_i^{k-1}(\ell+1)} \mid D^{k-1}\}. \end{aligned}$$

Thus, we get $\mathbb{E}\{e^{\lambda(M_i^{k-1}(\ell) + \sum_{s=1}^t n_{is}^k(q))} \mid D^{k-1} = q\} \leq e^{\frac{2}{\epsilon^6} \lambda^2} \mathbb{E}\{e^{\lambda(M_i^{k-1}(\ell+1) + \sum_{s=1}^t n_{is}^k(q))} \mid D^{k-1} = q\} \leq e^{\frac{2}{\epsilon^6} \lambda^2} e^{\frac{2}{\epsilon^6} (k-2-\ell)\lambda^2 + \gamma c_i \lambda}$, where the last inequality is by the induction assumption. \blacksquare

Note that $\{M_i^k(\ell) : \ell = 1, \dots, k\}$ is a martingale adapted to $\{\mathbf{D}^{[\ell, k]} : \ell = 1, \dots, k\}$ in the sense that $\mathbb{E}\{M_i^k(\ell) \mid \mathbf{D}^{[\ell+1, k]}\} = \mathbb{E}\{\mathbb{E}\{V_i^k \mid \mathbf{D}^{[\ell, k]}\} \mid \mathbf{D}^{[\ell+1, k]}\} = \mathbb{E}\{V_i^k \mid \mathbf{D}^{[\ell+1, k]}\} = M_i^k(\ell+1)$.

Performance Guarantee for the Approximate Policy:

In the next lemma, we use the moment generating function bounds given in Lemmas 5.1 and 5.2 to lower bound the availability probability on the right side of (4).

Lemma 5.3 (Availability Probability Bound) *Letting $U_i^0 = 0$, for all $k \in \mathcal{K}$, $i \in \mathcal{L}$, $t, q \in \mathcal{T}$ and $\lambda \in [0, 1]$, we have*

$$\mathbb{P}\left\{U_i^{k-1} + \sum_{s=1}^t N_{is}^k(q) \geq c_i \mid D^{k-1} = q\right\} \leq e^{(c_i + \frac{3}{\epsilon^6}(k-1))\lambda^2 - (1-\gamma)c_i \lambda}.$$

Proof: By the discussion just after the definition of $M_i^k(\ell)$, we have $V_i^{k-1} = M_i^{k-1}(1)$ with probability one, so $\mathbb{E}\{e^{\lambda(V_i^{k-1} + \sum_{s \in \mathcal{T}} n_{is}^k(q))} \mid D^{k-1} = q\} = \mathbb{E}\{e^{\lambda(M_i^{k-1}(1) + \sum_{s \in \mathcal{T}} n_{is}^k(q))} \mid D^{k-1} = q\} \leq e^{\frac{2}{\epsilon^6} (k-2)\lambda^2 + \gamma c_i \lambda}$, where the last inequality uses Lemma 5.2. On the other hand, by the first constraint in (2), we have $\sum_{s=1}^t n_{is}^k(q) = \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} \gamma \bar{x}_{js}^k(q) \leq \gamma c_i$. By a simple lemma, given as Lemma J.4 in Appendix J,

if the Bernoulli random variable Z has mean μ , then $\mathbb{E}\{e^{\lambda(Z-\mu)}\} \leq e^{\mu\lambda^2}$ for all $\lambda \in [0, 1]$, so because $n_{is}^k(q) = \mathbb{E}\{N_{is}^k(q)\}$, we get $\mathbb{E}\{e^{\lambda \sum_{s=1}^t (N_{is}^k(q) - n_{is}^k(q))}\} \leq e^{\lambda^2 \sum_{s=1}^t n_{is}^k(q)} \leq e^{\gamma c_i \lambda^2}$. We get

$$\begin{aligned} \mathbb{P}\left\{U_i^{k-1} + \sum_{s=1}^t N_{is}^k(q) \geq c_i \mid D^{k-1} = q\right\} &= \mathbb{P}\{e^{\lambda(U_i^{k-1} + \sum_{s=1}^t N_{is}^k(q))} \geq e^{\lambda c_i} \mid D^{k-1} = q\} \\ &\stackrel{(a)}{\leq} \frac{1}{e^{\lambda c_i}} \mathbb{E}\{e^{\lambda(U_i^{k-1} + \sum_{s=1}^t N_{is}^k(q))} \mid D^{k-1} = q\} \\ &\stackrel{(b)}{=} \frac{1}{e^{\lambda c_i}} \mathbb{E}\{e^{\lambda \sum_{s=1}^t (N_{is}^k(q) - n_{is}^k(q))}\} \mathbb{E}\{e^{\lambda(U_i^{k-1} + \sum_{s=1}^t n_{is}^k(q))} \mid D^{k-1} = q\} \\ &\stackrel{(c)}{\leq} \frac{1}{e^{\lambda c_i}} \mathbb{E}\{e^{\lambda \sum_{s=1}^t (N_{is}^k(q) - n_{is}^k(q))}\} e^{\frac{1}{2\epsilon}(k-1)\lambda^2} \mathbb{E}\{e^{\lambda(U_i^{k-1} + \sum_{s=1}^t n_{is}^k(q))} \mid D^{k-1} = q\} \\ &\stackrel{(d)}{\leq} \frac{1}{e^{\lambda c_i}} e^{\gamma c_i \lambda^2} e^{\frac{1}{2\epsilon}(k-1)\lambda^2} e^{\frac{2}{\epsilon^6}(k-2)\lambda^2 + \gamma c_i \lambda} \\ &= e^{(\gamma c_i + \frac{1}{2\epsilon}(k-1) + \frac{2}{\epsilon^6}(k-2))\lambda^2 - (1-\gamma)c_i \lambda} \stackrel{(e)}{\leq} e^{(c_i + \frac{3}{\epsilon^6}(k-1))\lambda^2 - (1-\gamma)c_i \lambda}, \end{aligned}$$

where (a) uses the Markov inequality, (b) holds because $N_{is}^k(q)$ is independent of D^{k-1} , (c) is by Lemma 5.1, (d) uses the two inequalities at the beginning of the proof and (e) uses $\epsilon \leq 1$. \blacksquare

Using specific values for γ and λ in Lemma 5.3, we will bound the availability probabilities. Using this bound in (3) will yield the performance guarantee in Theorem 4.1.

Proof of Theorem 4.1:

We use Lemma 5.3 with specific values of γ and λ . Letting $\delta = \frac{1}{\epsilon^6}$ for notational brevity, fix $\bar{\gamma} = 1 - \frac{\sqrt{4(c_{\min} + 3\delta(K-1)) \log c_{\min}}}{c_{\min}}$ and $\bar{\lambda} = \frac{(1-\bar{\gamma})c_i}{2(c_i + 3\delta(K-1))}$. For these values of $\bar{\gamma}$ and $\bar{\lambda}$, we have

$$(c_i + 3\delta(K-1))\bar{\lambda}^2 - (1-\bar{\gamma})c_i\bar{\lambda} \stackrel{(a)}{=} -\frac{[(1-\bar{\gamma})c_i]^2}{4(c_i + 3\delta(K-1))} \stackrel{(b)}{\leq} -\frac{[(1-\bar{\gamma})c_{\min}]^2}{4(c_{\min} + 3\delta(K-1))} \stackrel{(c)}{=} -\log c_{\min}, \quad (6)$$

where (a) follows by direct computation with the specific value of $\bar{\lambda}$, (b) holds because we can check the first derivative to verify that $\frac{[(1-\bar{\gamma})x]^2}{4(x+3\delta(K-1))}$ is increasing in x for $x \geq 0$ and (c) follows by noting that the value of $\bar{\gamma}$ satisfies $(1-\bar{\gamma})^2 c_{\min}^2 = 4(c_{\min} + 3\delta(K-1)) \log c_{\min}$. Without loss of generality, we can assume that $c_{\min} > 4\sqrt{(c_{\min} + \delta(K-1)) \log c_{\min}}$. Otherwise, the second term in the max operator in the theorem becomes a negative number and $\frac{\text{APX}}{\text{ZLP}}$ is trivially lower bounded by a negative number, so the result immediately holds. Therefore, we have $c_{\min} > 4\sqrt{(c_{\min} + \delta(K-1)) \log c_{\min}} \geq \sqrt{4(c_{\min} + 3\delta(K-1)) \log c_{\min}}$. In this case, our choice of $\bar{\gamma}$ satisfies $\bar{\gamma} \in [0, 1]$. If $\bar{\gamma} \in [0, 1]$, then our choice of $\bar{\lambda}$ satisfies $\bar{\lambda} \in [0, 1]$ as well. Therefore, we can use Lemma 5.3 with $\gamma = \bar{\gamma}$ and $\lambda = \bar{\lambda}$, so noting (6), we obtain

$$\mathbb{P}\left\{U_i^{k-1} + \sum_{s=1}^t N_{is}^k(q) \geq c_i \mid D^{k-1} = q\right\} \leq e^{(c_i + \frac{3}{\epsilon^6}(K-1))\bar{\lambda}^2 - (1-\bar{\gamma})c_i\bar{\lambda}} \leq e^{-\log c_{\min}} = \frac{1}{c_{\min}}.$$

By the definition of U_i^{k-1} , the probability on the left side above is the same as the probability on the right side of (4). Using the inequality above on the right side of (4), because $|\mathcal{L}_j| \leq L$, we

get $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} \geq 1 - \frac{L}{c_{\min}}$. By the discussion just after (3), if $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} \geq \alpha$, then $\frac{\text{APX}}{\bar{Z}_{\text{LP}}} \geq \gamma \alpha$. Thus, using the specific value of $\bar{\gamma}$, we get

$$\frac{\text{APX}}{\bar{Z}_{\text{LP}}} \geq \left(1 - \frac{\sqrt{4(c_{\min} + 3\delta(K-1)) \log c_{\min}}}{c_{\min}}\right) \left(1 - \frac{L}{c_{\min}}\right).$$

The result follows by noting that the right side of the chain of inequalities above is lower bounded by $1 - 4 \frac{\sqrt{(c_{\min} + \delta(K-1)) \log c_{\min}}}{c_{\min}} - \frac{L}{c_{\min}}$, as well as using the fact that $\frac{\text{APX}}{\text{OPT}} \geq \frac{\text{APX}}{\bar{Z}_{\text{LP}}}$ by Theorem 3.1. ■

In Appendix K, we unpack the proof of Theorem 4.1 further by elaborating on the tradeoffs that lead to the specific choices for the values of γ and λ in the proof.

6. Extensions, Numerical Performance and Research Directions

We can study performance guarantees when the demand random variables D^1, \dots, D^K in different stages are independent. In Appendix L, assuming that the demand random variable in each stage is sub-Gaussian with variance proxy σ^2 and independent of the demand random variables in other stages, we give an approximate policy with a performance guarantee of $1 - O\left(\frac{\sqrt{(c_{\min} + K\sigma^2) \log c_{\min}}}{c_{\min}}\right)$. In Appendix M, we show that our fluid approximation naturally extends to the case where we make assortment offering or pricing decisions. In Appendix N, we give computational experiments to demonstrate that policies driven by the correct fluid approximation can make a dramatic impact in practice. Our work opens up several research directions. It would be useful to test our model in practical revenue management systems. The additional input requirement of our model is the distribution of the customer arrivals in the next stage conditional on the number of customer arrivals in the current stage. Assuming that each stage corresponds to a week, one approach to estimate this input requirement is that we start by unconstraining the customer arrivals so that we obtain the raw data that corresponds to the volume of the customer arrivals in each week; see Weatherford and Polt (2002), Queenan et al. (2007) and Kourentzes et al. (2019). Obtaining such unconstrained customer arrival volumes is often done in revenue management systems. Next, we can fit an autoregressive process to the customer arrival data; see Box et al. (2015). Lastly, there is work on capturing the evolution of an autoregressive process through a Markov chain, allowing us to capture the distribution of the demand at the next stage as a function of the demand in the current stage; see Tauchen (1986), Burnside (2001) and Fella et al. (2019). This approach is one possibility for estimating the key input requirement of our model and it is useful to test such approaches in practice. Another research direction is that the demand in a stage depends on the demand in the previous stage in our model. One can focus on more complicated dependence structures for the demand.

Acknowledgements: The authors thank the department editor, associate editor and two referees whose comments significantly improved the depth and exposition in the paper.

References

- Adelman, D. 2007. Dynamic bid-prices in revenue management. *Operations Research* **55**(4) 647–661.
- Adelman, D., A. J. Mersereau. 2008. Relaxations of weakly coupled stochastic dynamic programs. *Operations Research* **56**(3) 712–727.
- Aouad, A., W. Ma. 2022. A nonparametric framework for online stochastic matching with correlated arrivals. Tech. rep., London Business School, London, UK.
- Baek, J., W. Ma. 2022. Technical note – Bifurcating constraints to improve approximation ratios for network revenue management with reusable resources. *Operations Research* **70**(4) 2226–2236.
- Bai, Y., O. El Housni, B. Jin, P. Rusmevichientong, H. Topaloglu, D. P. Williamson. 2023. Fluid approximations for revenue management under high-variance demand. *Management Science* **69**(7) 4016–4026.
- Bai, Y., O. El Housni, P. Rusmevichientong, H. Topaloglu. 2022. Coordinated inventory stocking and assortment personalization. Tech. rep., Cornell University, New York, NY.
- Balseiro, S. R., O. Besbes, D. Pizarro. 2023. Survey of dynamic resource constrained reward collection problems: Unified model and analysis. *Operations Research* (to appear).
- Box, G. E. P, G. M. Jenkins, G. C. Reinsel, G. M. Ljung. 2015. *Time Series Analysis: Forecasting and Control*. John Wiley & Sons, Hoboken, NJ.
- Burnside, C. 2001. Discrete State-Space Methods for the Study of Dynamic Economies. *Computational Methods for the Study of Dynamic Economies*. Oxford University Press, New York, NY.
- Chen, L. H. Y, L. Goldstein, Q. M. Shao. 2010. *Normal Approximation by Stein's Method*. Probability and Its Applications, Springer, Berlin, Germany.
- Dubhashi, D. P., A. Panconesi. 2009. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, Cambridge, UK.
- Fella, G., G. Gallipoli, J. Pan. 2019. Markov-chain approximations for life-cycle models. *Review of Economic Dynamics* **34** 183–201.
- Feng, Y., R. Niazadeh, A. Saberi. 2022. Near-optimal Bayesian online assortment of reusable resources. Tech. rep., University of Chicago, Chicago, IL.
- Gallego, G., G. Iyengar, R. Phillips, A. Dubey. 2004. Managing flexible products on a network. CORC Technical Report TR-2004-01.
- Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* **40**(8) 999–1020.
- Gallego, G., G. van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research* **45**(1) 24–41.
- Jasin, S., S. Kumar. 2012. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research* **37**(2) 313–345.
- Jiang, J. 2023. Constant approximation for network revenue management with Markovian-correlated customer arrivals. Tech. rep., Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong.
- Kourentzes, N., F. Li, A. K. Strauss. 2019. Unconstraining methods for revenue management systems under small demand. *Journal of Revenue and Pricing Management* **18** 27–41.
- Kunnumkal, S., K. Talluri, H. Topaloglu. 2012. A randomized linear programming method for network revenue management with product-specific no-shows. *Transportation Science* **46**(1) 90–108.
- Levin, D. A., Y. Peres. 2017. *Markov Chains and Mixing Times*. American Mathematical Society, Providence, RI.
- Liu, Q., G. J. van Ryzin. 2008. On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management* **10**(2) 288–310.
- Ma, W., D. Simchi-Levi, J. Zhao. 2021. Dynamic pricing (and assortment) under a static calendar. *Management Science* **67**(4) 2292–2313.
- Ma, Y., P. Rusmevichientong, M. Sumida, H. Topaloglu. 2020. An approximation algorithm for network revenue management under nonstationary arrivals. *Operations Research* **68**(3) 834–855.
- Queenan, C. C., M. Ferguson, J. Higbie, R. Kapoor. 2007. A comparison of unconstraining methods to improve revenue management systems. *Production and Operations Management* **16**(6) 729–746.
- Rusmevichientong, P., M. Sumida, H. Topaloglu. 2020. Dynamic assortment optimization for reusable products with random usage durations. *Management Science* **66**(7) 2820–2844.
- Talluri, K., G. van Ryzin. 1998. An analysis of bid-price controls for network revenue management. *Management Science* **44**(11) 1577–1593.
- Talluri, K., G. van Ryzin. 1999. A randomized linear programming method for computing network bid prices. *Transportation Science* **33**(2) 207–216.
- Tauchen, G. 1986. Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters* **20**(2) 177–181.
- Wainwright, M. J. 2019. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Cambridge, UK.
- Weatherford, L. R., S. Polt. 2002. Better unconstraining of airline demand data in revenue management systems for improved forecast accuracy and greater revenues. *Journal of Revenue and Pricing Management* **1**(3) 234–254.

Electronic Supplement:
Technical Note – Revenue Management with Calendar-Aware and
Dependent Demands: Asymptotically Tight Fluid Approximations

Weiyuan Li, Paat Rusmevichientong, Huseyin Topaloglu
October 19, 2024

Online Appendix:

Technical Note – Revenue Management with Calendar-Aware and Dependent Demands: Asymptotically Tight Fluid Approximations

Weiyuan Li, Paat Rusmevichientong, Huseyin Topaloglu

Appendix A: Matching a Given Total Demand Distribution

Consider a target distribution for the total demand over the selling horizon with a support bounded by KT . We show that we can calibrate our demand model such that the distribution of the total demand implied by our demand model matches the target distribution. In our demand model, there are K stages and the demand in each stage takes values between one and T , so the total demand takes values in $\{K, \dots, KT\}$. The assumption that the demand in each stage is at least one is without loss of generality. We can add a dummy time period to each stage such that the probability that there is a request for any product at the dummy time period is zero, in which case, there is at least one time period in each stage, but we may not have any demand in a particular stage. Nevertheless, for consistency with our demand model, we consider an arbitrary total demand random variable D^{total} that takes values in $\{K, \dots, KT\}$. We show that we can calibrate our demand model such that the distribution of the total demand implied by our demand model matches the distribution of D^{total} . Assuming that we are given the random variable D^{total} for the total demand, for stage k , we define the probability mass function over $\{1, \dots, T\}$ as

$$f^k(\ell) = \begin{cases} \mathbb{P}\{D^{\text{total}} = (k-1)(T-1) + K + \ell - 1 \mid D^{\text{total}} \geq (k-1)(T-1) + K\} & \text{if } \ell = 1, \dots, T-1 \\ \mathbb{P}\{D^{\text{total}} \geq (k-1)(T-1) + K + T - 1 \mid D^{\text{total}} \geq (k-1)(T-1) + K\} & \text{if } \ell = T. \end{cases} \quad (7)$$

Note that we have $\sum_{\ell=1}^T f^k(\ell) = 1$, so f^k is indeed a probability mass function. The way we use the probability mass function shortly becomes clear, but to provide intuition into this probability mass function, note that the random variable $D^{\text{total}} - K$ takes values in $\{0, \dots, K(T-1)\}$. In our model, on the other hand, the demand in each stage takes at least a value of one. In addition to this value of one, the demand in each stage can contribute at most $T-1$ units to the total demand. In this case, we observe that if the condition $D^{\text{total}} \geq (k-1)(T-1) + K$ in the conditional probability in (7) holds, then we have $D^{\text{total}} - K \geq (k-1)(T-1)$, which implies that the total demand surely exceeds the demands in the first $k-1$ stages of our model. In the rest of this section, we will construct the random variables D^1, \dots, D^K that satisfy the following three properties. First, each of the random variables D^1, \dots, D^K takes values in $\{1, \dots, T\}$. Second, conditional on D^k , the random variable D^{k+1} is independent of the random variables D^1, \dots, D^{k-1} . Third, the random variable $D^1 + \dots + D^K$ has the same distribution as the random variable D^{total} . We accomplish this goal as follows. We use Z^1, \dots, Z^K to denote a sequence of independent random variables, where Z^k has

the probability mass function f^k . Using the random variables Z^1, \dots, Z^K , we define the random variables D^1, \dots, D^K recursively as

$$D^k = \begin{cases} 1 & \text{if } D^{k-1} \neq T \\ Z^k & \text{if } D^{k-1} = T, \end{cases} \quad (8)$$

where we fix $D^0 = T$. Note that D^k takes values in $\{1, \dots, T\}$. Also, D^k depends only on D^{k-1} and Z^k . Thus, since the random variables Z^1, \dots, Z^K are independent of each other, given D^{k-1} , the random variable D^k is independent of D^1, \dots, D^{k-2} . In the next theorem, we show that the random variable $D^1 + \dots + D^K$ has the same distribution as D^{total} . Before we state this theorem, by (8), note that only one of the random variables D^1, \dots, D^K can take a value other than one or T . Also, if the random variable D^k takes a value other than one or T , then the random variables D^1, \dots, D^{k-1} all take value T and the random variables D^{k+1}, \dots, D^K all take value one.

Theorem A.1 *Letting the random variables D^1, \dots, D^K be recursively defined as in (8), the random $D^1 + \dots + D^K$ has the same distribution as the random variable D^{total} .*

Proof: Fixing some $\ell \in \{K, \dots, KT\}$, we show that $\mathbb{P}\{D^1 + \dots + D^K = \ell\} = \mathbb{P}\{D^{\text{total}} = \ell\}$. Using $\lceil \cdot \rceil$ to denote the round up function, set $r = \lceil \frac{\ell - K}{T-1} \rceil$, which takes values in $\{0, \dots, K\}$. By the discussion just before the theorem, only one of the random variables in the sequence D^1, \dots, D^K can take a value other than one or T . If D^k is such a random variable, then all of the random variables D^1, \dots, D^{k-1} take value T and all of the random variables D^{k+1}, \dots, D^K take value one. In this case, if $D^1 + \dots + D^K = \ell$, then we must have $D^1 = T, \dots, D^{r-1} = T, D^{r+1} = 1, \dots, D^K = 1$ and $D^r = \ell - (r-1)T - (K-r)$. Considering the case $\ell - (r-1)T - (K-r) \neq T$, we have

$$\begin{aligned} & \mathbb{P}\{D^1 + \dots + D^K = \ell\} \\ &= \mathbb{P}\{D^1 = T, \dots, D^{r-1} = T, D^r = \ell - (r-1)T - (K-r), D^{r+1} = 1, \dots, D^K = 1\} \\ &\stackrel{(a)}{=} \mathbb{P}\{D^1 = T\} \times \mathbb{P}\{D^2 = T \mid D^1 = T\} \times \dots \times \mathbb{P}\{D^r = \ell - (r-1)T - (K-r) \mid D^{r-1} = T\} \\ &\quad \times \mathbb{P}\{D^{r+1} = 1 \mid D^r = \ell - (r-1)T - (K-r)\} \\ &\quad \times \mathbb{P}\{D^{r+2} = 1 \mid D^{r+1} = 1\} \times \dots \times \mathbb{P}\{D^K = 1 \mid D^{K-1} = 1\} \\ &\stackrel{(b)}{=} \mathbb{P}\{Z^1 = T\} \times \mathbb{P}\{Z^2 = T\} \times \dots \times \mathbb{P}\{Z^r = \ell - (r-1)T - (K-r)\} \\ &\stackrel{(c)}{=} \mathbb{P}\{D^{\text{total}} \geq T-1 + K \mid D^{\text{total}} \geq K\} \times \mathbb{P}\{D^{\text{total}} \geq 2(T-1) + K \mid D^{\text{total}} \geq T-1 + K\} \\ &\quad \times \dots \times \mathbb{P}\{D^{\text{total}} = (r-1)(T-1) + K + \ell - (r-1)T - (K-r) - 1 \mid D^{\text{total}} \geq (r-1)(T-1) + K\} \\ &\stackrel{(d)}{=} \mathbb{P}\{D^{\text{total}} = \ell, D^{\text{total}} \geq (r-1)(T-1) + K \mid D^{\text{total}} \geq K\} \stackrel{(e)}{=} \mathbb{P}\{D^{\text{total}} = \ell\}, \end{aligned}$$

where (a) holds because conditional on D^k , the random variable D^{k+1} is independent of the random variables D^1, \dots, D^{k-1} , (b) uses (8) along with the fact that $\ell - (r-1)T - (K-r) \neq T$, (c) follows

from (7), (d) uses the Bayes rule and (e) holds because the support of D^{total} is lower bounded by K , as well as the fact that $r \leq \frac{\ell-K}{T-1}$, which implies that $\ell \geq r(T-1) + K$. Considering the case $\ell - (r-1)T - (K-r) = T$, we can follow an argument similar to the one in this paragraph to show that we have $\mathbb{P}\{D^1 + \dots + D^K = \ell\} = \mathbb{P}\{D^{\text{total}} = \ell\}$ in this case as well. ■

Appendix B: Proof of Theorem 3.1

We relax the constraint $a_{ij} u_j \leq y_i$ at time period t in stage k in (1) using the Lagrange multiplier $\mu_{it}^k(q)$. Letting $\boldsymbol{\mu} = (\mu_{it}^k(q) : i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K})$, we obtain the relaxed dynamic program

$$\begin{aligned} \tilde{J}_t^k(\mathbf{y}, q; \boldsymbol{\mu}) &= \max_{\mathbf{u} \in \{0,1\}^{|\mathcal{J}|}} \left\{ \sum_{j \in \mathcal{T}} \lambda_{jt}^k \left\{ f_j u_j + \theta_t^k(q) \tilde{J}_{t+1}^k(\mathbf{y} - \mathbf{a}_j u_j, q; \boldsymbol{\mu}) \right. \right. \\ &\quad \left. \left. + (1 - \theta_t^k(q)) \tilde{J}_1^{k+1}(\mathbf{y} - \mathbf{a}_j u_j, t; \boldsymbol{\mu}) \right\} + \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{J}} \lambda_{jt}^k \mu_{it}^k(q) [y_i - a_{ij} u_j] \right\} \\ &= \max_{\mathbf{u} \in \{0,1\}^{|\mathcal{J}|}} \left\{ \sum_{j \in \mathcal{T}} \lambda_{jt}^k \left\{ \left[f_j - \sum_{i \in \mathcal{L}} a_{ij} \mu_{it}^k(q) \right] u_j + \theta_t^k(q) \tilde{J}_{t+1}^k(\mathbf{y} - \mathbf{a}_j u_j, q; \boldsymbol{\mu}) \right. \right. \\ &\quad \left. \left. + (1 - \theta_t^k(q)) \tilde{J}_1^{k+1}(\mathbf{y} - \mathbf{a}_j u_j, t; \boldsymbol{\mu}) \right\} \right\} + \sum_{i \in \mathcal{L}} \mu_{it}^k(q) y_i, \end{aligned} \quad (9)$$

with the boundary condition that $\tilde{J}_1^{K+1} = 0$. Note that the value functions of the relaxed dynamic program depend on the choice of the Lagrange multipliers. In the first equality above, we scale the Lagrange multiplier $\mu_{it}^k(q)$ with λ_{jt}^k for notational uniformity. The second equality follows by arranging the terms and using the fact that $\sum_{j \in \mathcal{J}} \lambda_{jt}^k = 1$. If the Lagrange multipliers are non-negative, then the value functions from the relaxed dynamic program in (9) are upper bounds on the value functions from the dynamic program in (1). We do not show this result. This result is considered standard and analogues of this result have been shown in other settings; see Proposition 2 in Adelman and Mersereau (2008). Therefore, we have $\tilde{J}_t^k(\mathbf{y}, q; \boldsymbol{\mu}) \geq J_t^k(\mathbf{y}, q)$ for all $\mathbf{y} \in \mathbb{Z}_+^{|\mathcal{L}|}$, $q \in \mathcal{T}$ as long as $\boldsymbol{\mu} \in \mathbb{R}_+^{|\mathcal{L}|T^2K}$. We can solve the problem $\min_{\boldsymbol{\mu} \in \mathbb{R}_+^{|\mathcal{L}|T^2K}} \tilde{J}_1^1(\mathbf{c}, \hat{D}^0; \boldsymbol{\mu})$ to obtain an upper bound on the optimal total expected revenue. One of the useful features of the relaxed dynamic program is that the value functions computed through this dynamic program are linear in the remaining capacities. In the next lemma, we show that $\tilde{J}_t^k(\mathbf{y}, q; \boldsymbol{\mu}) = \sum_{i \in \mathcal{L}} \alpha_{it}^k(q; \boldsymbol{\mu}) y_i + \beta_t^k(q; \boldsymbol{\mu})$, where the slope $\alpha_{it}^k(q; \boldsymbol{\mu})$ and the intercept $\beta_t^k(q; \boldsymbol{\mu})$ are recursively computed as

$$\begin{aligned} \alpha_{it}^k(q; \boldsymbol{\mu}) &= \mu_{it}^k(q) + \theta_t^k(q) \alpha_{i,t+1}^k(q; \boldsymbol{\mu}) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t; \boldsymbol{\mu}) \\ \beta_t^k(q; \boldsymbol{\mu}) &= \sum_{j \in \mathcal{J}} \lambda_{jt}^k \left[f_j - \sum_{i \in \mathcal{L}} a_{ij} \alpha_{it}^k(q; \boldsymbol{\mu}) \right]^+ + \theta_t^k(q) \beta_{t+1}^k(q; \boldsymbol{\mu}) + (1 - \theta_t^k(q)) \beta_1^{k+1}(t; \boldsymbol{\mu}), \end{aligned} \quad (10)$$

with the boundary condition that $\alpha_{i1}^{K+1} = 0$ and $\beta_1^{K+1} = 0$. The linear form of the value functions from the relaxed dynamic program will be useful to show Theorem 3.1.

Lemma B.1 *Letting $\alpha_{it}^k(q; \boldsymbol{\mu})$ and $\beta_t^k(q; \boldsymbol{\mu})$ be as in (10), the value functions computed through the dynamic program in (9) satisfy $\tilde{J}_t^k(\mathbf{y}, q; \boldsymbol{\mu}) = \sum_{i \in \mathcal{L}} \alpha_{it}^k(q; \boldsymbol{\mu}) y_i + \beta_t^k(q; \boldsymbol{\mu})$ for all $t \in \mathcal{T}$ and $k \in \mathcal{K}$.*

Proof: We show the result by using induction over the time periods. At the last time period in the last stage, by (9), we have $J_T^K(\mathbf{y}, q; \boldsymbol{\mu}) = \sum_{j \in \mathcal{T}} \lambda_{jT}^K [f_j - \sum_{i \in \mathcal{L}} a_{ij} \mu_{iT}^K(q)]^+ + \sum_{i \in \mathcal{L}} \mu_{iT}^K(q) y_i = \beta_T^K(q; \boldsymbol{\mu}) + \sum_{i \in \mathcal{L}} \alpha_{iT}^K(q; \boldsymbol{\mu}) y_i$, where the last equality uses (10). Therefore, the result holds at the last time period in the last stage. Assuming that the result holds at all time periods after time period t in stage k , we show that the result holds at time period t in stage k as well. Using the induction assumption on the right side of (9), we have

$$\begin{aligned} \tilde{J}_t^k(\mathbf{y}, q; \boldsymbol{\mu}) &= \max_{\mathbf{u} \in \{0,1\}^{|\mathcal{J}|}} \left\{ \sum_{j \in \mathcal{T}} \lambda_{jt}^k \left\{ \left[f_j - \sum_{i \in \mathcal{L}} a_{ij} \mu_{it}^k(q) \right] u_j + \theta_t^k(q) \left[\beta_{t+1}^k(q; \boldsymbol{\mu}) + \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q; \boldsymbol{\mu}) (y_i - a_{ij} u_j) \right] \right. \right. \\ &\quad \left. \left. + (1 - \theta_t^k(q)) \left[\beta_1^{k+1}(t; \boldsymbol{\mu}) + \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t; \boldsymbol{\mu}) (y_i - a_{ij} u_j) \right] \right\} \right\} + \sum_{i \in \mathcal{L}} \mu_{it}^k(q) y_i \\ &\stackrel{(a)}{=} \max_{\mathbf{u} \in \{0,1\}^{|\mathcal{J}|}} \left\{ \sum_{j \in \mathcal{T}} \lambda_{jt}^k \left[f_j - \sum_{i \in \mathcal{L}} a_{ij} \alpha_{it}^k(q; \boldsymbol{\mu}) \right] u_j \right\} \\ &\quad + \theta_t^k(q) \beta_{t+1}^k(q; \boldsymbol{\mu}) + (1 - \theta_t^k(q)) \beta_1^{k+1}(t; \boldsymbol{\mu}) + \sum_{i \in \mathcal{L}} \alpha_{it}^k(q; \boldsymbol{\mu}) y_i \\ &= \sum_{j \in \mathcal{T}} \lambda_{jt}^k \left[f_j - \sum_{i \in \mathcal{L}} a_{ij} \alpha_{it}^k(q; \boldsymbol{\mu}) \right]^+ + \theta_t^k(q) \beta_{t+1}^k(q; \boldsymbol{\mu}) + (1 - \theta_t^k(q)) \beta_1^{k+1}(t; \boldsymbol{\mu}) + \sum_{i \in \mathcal{L}} \alpha_{it}^k(q; \boldsymbol{\mu}) y_i \\ &\stackrel{(b)}{=} \beta_t^k(q; \boldsymbol{\mu}) + \sum_{i \in \mathcal{L}} \alpha_{it}^k(q; \boldsymbol{\mu}) y_i, \end{aligned}$$

where (a) follows by arranging the terms and using the definition of $\alpha_{it}^k(q; \boldsymbol{\mu})$, as well as noting the fact that $\sum_{j \in \mathcal{J}} \lambda_{jt}^k = 1$, whereas (b) uses the definition of $\beta_t^k(q; \boldsymbol{\mu})$. \blacksquare

By the lemma above, we have $\tilde{J}_1^k(\mathbf{c}, \hat{D}^0; \boldsymbol{\mu}) = \sum_{i \in \mathcal{L}} \alpha_{i1}^k(\hat{D}^0; \boldsymbol{\mu}) c_i + \beta_1^k(\hat{D}^0; \boldsymbol{\mu})$. In this case, the problem $\min_{\boldsymbol{\mu} \in \mathbb{R}_+^{|\mathcal{L}|T^2K}} \tilde{J}_1^k(\mathbf{c}, \hat{D}^0; \boldsymbol{\mu})$ is equivalent to the linear program

$$\begin{aligned} \min_{(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\eta}) \in \mathbb{R}^{T^2K(|\mathcal{L}|+1)} \times \mathbb{R}_+^{T^2K(|\mathcal{L}|+|\mathcal{J}|)}} & \sum_{i \in \mathcal{L}} \alpha_{i1}^1(\hat{D}^0) c_i + \beta_1^1(\hat{D}^0) \tag{11} \\ \text{st} & \alpha_{it}^k(q) = \mu_{it}^k(q) + \theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t) \quad \forall i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K} \\ & \beta_t^k(q) = \sum_{j \in \mathcal{J}} \lambda_{jt}^k \eta_{jt}^k(q) + \theta_t^k(q) \beta_{t+1}^k(q) + (1 - \theta_t^k(q)) \beta_1^{k+1}(t) \quad \forall t, q \in \mathcal{T}, k \in \mathcal{K} \\ & \eta_{jt}^k(q) \geq f_j - \sum_{i \in \mathcal{L}} a_{ij} \alpha_{it}^k(q) \quad \forall j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K}, \end{aligned}$$

where we use the decision variables $\boldsymbol{\alpha} = (\alpha_{it}^k(q) : i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K})$, $\boldsymbol{\beta} = (\beta_t^k(q) : t \in \mathcal{T}, k \in \mathcal{K})$, $\boldsymbol{\mu} = (\mu_{it}^k(q) : i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K})$ and $\boldsymbol{\eta} = (\eta_{jt}^k(q) : j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$. We follow the convention

that $\alpha_{i1}^{K+1}(q) = 0$ and $\beta_1^{K+1}(q) = 0$ for all $i \in \mathcal{L}$ and $q \in \mathcal{T}$. In the linear program above, the first constraint computes the slopes of the value functions of the relaxed dynamic program. By the third constraint, noting the non-negativity constraints, we have $\eta_{jt}^k(q) = [f_j - \sum_{i \in \mathcal{L}} a_{ij} \alpha_{it}^k(q)]^+$ at an optimal solution to the linear program, in which case, the second constraint computes the intercepts of the value functions of the relaxed dynamic program. We write the objective function of problem (11) as $\sum_{q \in \mathcal{T}} \sum_{i \in \mathcal{L}} \mathbf{1}(\widehat{D}^0 = q) c_i \alpha_{i1}^1(q) + \sum_{q \in \mathcal{T}} \mathbf{1}(\widehat{D}^0 = q) \beta_1^1(q)$. We work with the dual of problem (11). We associate the dual variables $\mathbf{w} = (w_t^k(q) : t, q \in \mathcal{T}, k \in \mathcal{K})$ with the second constraint in (11). The decision variables $(\beta_t^k(q) : t, q \in \mathcal{T}, k \in \mathcal{K})$ appear only in the second constraint in problem (11), so in the dual of problem (11), the constraints associated with the decision variables $(\beta_t^k(q) : t, q \in \mathcal{T}, k \in \mathcal{K})$ are given by $w_1^1(q) = \mathbf{1}(\widehat{D}^0 = q)$ for all $q \in \mathcal{T}$, $w_1^k(q) = \sum_{p \in \mathcal{T}} (1 - \theta_q^{k-1}(p)) w_q^{k-1}(p)$ for all $q \in \mathcal{T}, k \in \mathcal{K} \setminus \{1\}$ and $w_t^k(q) = \theta_{t-1}^k(q) w_{t-1}^k(q)$ for all $t \in \mathcal{T} \setminus \{1\}, q \in \mathcal{T}, k \in \mathcal{K}$. We capture these constraints by defining the set

$$\mathcal{W} = \left\{ \mathbf{w} \in \mathbb{R}^{T^2 K} : w_t^k(q) = \theta_{t-1}^k(q) w_{t-1}^k(q) \quad \forall t \in \mathcal{T} \setminus \{1\}, q \in \mathcal{T}, k \in \mathcal{K}, \right. \\ \left. w_1^k(q) = \sum_{p \in \mathcal{T}} (1 - \theta_q^{k-1}(p)) w_q^{k-1}(p) \quad \forall q \in \mathcal{T}, k \in \mathcal{K} \setminus \{1\}, w_1^1(q) = \mathbf{1}(\widehat{D}^0 = q) \quad \forall q \in \mathcal{T} \right\}.$$

In the next lemma, we show that if we have $\mathbf{w} \in \mathcal{W}$, then the vector \mathbf{w} is closely related to the joint distribution of demands in a pair of successive stages.

Lemma B.2 *If $\mathbf{w} \in \mathcal{W}$, then we have $w_t^k(q) = \mathbb{P}\{D^k \geq t, D^{k-1} = q\}$ for all $t, q \in \mathcal{T}$ and $k \in \mathcal{K}$.*

Proof: We show the result by using induction over the time periods. At the first time period in the first stage, we have $w_1^1(q) = \mathbf{1}(\widehat{D}^0 = q) = \mathbb{P}\{D^0 = q\} = \mathbb{P}\{D^1 \geq 1, D^0 = q\}$, where the first equality holds by noting the third constraint in the definition of \mathcal{W} , the second equality follows by noting that D^0 is deterministically fixed at \widehat{D}^0 and the third equality holds because the support of D^1 is $\{1, \dots, T\}$. Assuming that the result holds at all time periods before time period t in stage k , we show that the result holds at time period t in stage k as well. If $t \neq 1$, then using the first constraint in the definition of \mathcal{W} , we have the chain of equalities $w_t^k(q) = \theta_{t-1}^k(q) w_{t-1}^k(q) = \mathbb{P}\{D^k \geq t \mid D^k \geq t-1, D^{k-1} = q\} \mathbb{P}\{D^k \geq t-1, D^{k-1} = q\} = \mathbb{P}\{D^k \geq t, D^{k-1} = q\}$, where the second equality is by the definition of $\theta_{t-1}^k(q)$ and the induction assumption. Similarly, if $t = 1$, then using the second constraint in the definition of \mathcal{W} , we have $w_1^k(q) = \sum_{p \in \mathcal{T}} (1 - \theta_q^{k-1}(p)) w_q^{k-1}(p) = \sum_{p \in \mathcal{T}} \mathbb{P}\{D^{k-1} = q \mid D^{k-1} \geq q, D^{k-2} = p\} \mathbb{P}\{D^{k-1} \geq q, D^{k-2} = p\}$, but the last sum expression is equal to $\mathbb{P}\{D^{k-1} = q\}$, so $w_1^k(q) = \mathbb{P}\{D^{k-1} = q\} = \mathbb{P}\{D^k \geq 1, D^{k-1} = q\}$. ■

By the lemma above, there exists a single element in \mathcal{W} . To write the dual of problem (11), we associate the dual variables $\mathbf{y} = (y_{it}^k(q) : i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K})$, $\mathbf{w} = (w_t^k(q) : t, q \in \mathcal{T}, k \in \mathcal{K})$ and

$\mathbf{u} = (u_{jt}^k(q) : j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$ with the first, second and third constraints, respectively, in problem (11). In this case, the dual of problem (11) is given by

$$\begin{aligned}
& \max_{(\mathbf{y}, \mathbf{u}, \mathbf{w}) \in \mathbb{R}_+^{T^2 K(|\mathcal{L}| + |\mathcal{J}|)} \times \mathcal{W}} \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{q \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j u_{jt}^k(q) & (12) \\
& \text{st } y_{it}^k(q) + \sum_{j \in \mathcal{J}} a_{ij} u_{jt}^k(q) = \theta_{t-1}^k(q) y_{i,t-1}^k(q) \quad \forall i \in \mathcal{L}, t \in \mathcal{T} \setminus \{1\}, q \in \mathcal{T}, k \in \mathcal{K} \\
& y_{i1}^k(q) + \sum_{j \in \mathcal{J}} a_{ij} u_{j1}^k(q) = \sum_{p \in \mathcal{T}} (1 - \theta_q^{k-1}(p)) y_{iq}^{k-1}(p) \quad \forall i \in \mathcal{L}, q \in \mathcal{T}, k \in \mathcal{K} \setminus \{1\} \\
& y_{i1}^1(q) + \sum_{j \in \mathcal{J}} a_{ij} u_{j1}^1(q) = \mathbf{1}(\widehat{D}^0 = q) c_i \quad \forall i \in \mathcal{L}, q \in \mathcal{T} \\
& u_{jt}^k(q) \leq \lambda_{jt}^k w_t^k(q) \quad \forall j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K},
\end{aligned}$$

where the constraints above are associated with the decision variables $\boldsymbol{\alpha}$ and $\boldsymbol{\eta}$ in (11). The constraint for the decision variables $\boldsymbol{\mu}$ translates into the non-negativity constraint for \mathbf{y} .

In (12), we capture the constraint associated with the decision variables $\boldsymbol{\beta}$ as $\mathbf{w} \in \mathcal{W}$. In the next lemma, we give an equality that is satisfied by all feasible solutions to problem (12).

Lemma B.3 *Letting $(\mathbf{y}, \mathbf{u}, \mathbf{w})$ be a feasible solution to the linear program in (12), for all $i \in \mathcal{L}$, $t, q \in \mathcal{T}$ and $k \in \mathcal{K}$, we have*

$$\begin{aligned}
& \mathbb{P}\{D^k \geq t, D^{k-1} = q\} c_i - \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^k \geq t, D^{k-1} = q \mid D^\ell \geq s, D^{\ell-1} = p\} u_{js}^\ell(p) \\
& \quad - \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^k \geq t \mid D^k \geq s, D^{k-1} = q\} u_{js}^k(q) = y_{it}^k(q).
\end{aligned}$$

Proof: We show the result by using induction over the time periods. At the first time period in the first stage, by the third constraint in (12), we have $y_{i1}^1(q) = \mathbf{1}(\widehat{D}^0 = q) c_i - \sum_{j \in \mathcal{J}} a_{ij} u_{j1}^1(q) = \mathbb{P}\{D^1 \geq 1, D^0 = q\} c_i - \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^1 \geq 1 \mid D^1 \geq 1, D^0 = q\} u_{j1}^1(q)$, where the last equality holds because D^0 is a deterministic quantity and the support of D^1 is $\{1, \dots, T\}$. Assuming that the result holds at all time periods up to and including time period t in stage k , we show that the result holds at the subsequent time period as well. Consider the case $t \neq T$. We will use three identities. First, for $\ell \leq k-1$, given D^{k-1}, D^k is independent of D^1, \dots, D^ℓ , in which case, we obtain

$$\begin{aligned}
& \theta_t^k(q) \mathbb{P}\{D^k \geq t, D^{k-1} = q \mid D^\ell \geq s, D^{\ell-1} = p\} \\
& = \mathbb{P}\{D^k \geq t+1 \mid D^k \geq t, D^{k-1} = q\} \mathbb{P}\{D^k \geq t, D^{k-1} = q \mid D^\ell \geq s, D^{\ell-1} = p\} \\
& = \mathbb{P}\{D^k \geq t+1, D^{k-1} = q \mid D^\ell \geq s, D^{\ell-1} = p\}.
\end{aligned}$$

Second, by the Bayes rule and definition of $\theta_t^k(q)$, we can show that $\theta_t^k(q) \mathbb{P}\{D^k \geq t, D^{k-1} = q\} = \mathbb{P}\{D^k \geq t+1, D^{k-1} = q\}$. Third, for $s \leq t$, we can, once more, use the Bayes rule and definition of

$\theta_t^k(q)$ to show that $\theta_t^k(q) \mathbb{P}\{D^k \geq t \mid D^k \geq s, D^{k-1} = q\} = \mathbb{P}\{D^k \geq t+1 \mid D^k \geq s, D^{k-1} = q\}$. Noting that the solution $(\mathbf{y}, \mathbf{u}, \mathbf{w})$ is feasible to (12), it satisfies the first constraint. Thus, we obtain

$$\begin{aligned}
y_{i,t+1}^k(q) &= \theta_t^k(q) y_{it}^k(q) - \sum_{j \in \mathcal{J}} a_{ij} u_{j,t+1}^k(q) \\
&\stackrel{(a)}{=} \theta_t^k(q) \mathbb{P}\{D^k \geq t, D^{k-1} = q\} c_i \\
&\quad - \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \theta_t^k(q) \mathbb{P}\{D^k \geq t, D^{k-1} = q \mid D^\ell \geq s, D^{\ell-1} = p\} u_{js}^\ell(p) \\
&\quad - \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} \theta_t^k(q) \mathbb{P}\{D^k \geq t \mid D^k \geq s, D^{k-1} = q\} u_{js}^k(q) - \sum_{j \in \mathcal{J}} a_{ij} u_{j,t+1}^k(q) \\
&\stackrel{(b)}{=} \mathbb{P}\{D^k \geq t+1, D^{k-1} = q\} c_i \\
&\quad - \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^k \geq t+1, D^{k-1} = q \mid D^\ell \geq s, D^{\ell-1} = p\} u_{js}^\ell(p) \\
&\quad - \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^k \geq t+1 \mid D^k \geq s, D^{k-1} = q\} u_{js}^k(q) - \sum_{j \in \mathcal{J}} a_{ij} u_{j,t+1}^k(q) \\
&\stackrel{(c)}{=} \mathbb{P}\{D^k \geq t+1, D^{k-1} = q\} c_i \\
&\quad - \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^k \geq t+1, D^{k-1} = q \mid D^\ell \geq s, D^{\ell-1} = p\} u_{js}^\ell(p) \\
&\quad - \sum_{s=1}^{t+1} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^k \geq t+1 \mid D^k \geq s, D^{k-1} = q\} u_{js}^k(q),
\end{aligned}$$

where (a) is by the induction assumption, (b) uses the three identities given earlier in the proof and (c) holds by noting that $\mathbb{P}\{D^k \geq t+1 \mid D^k \geq t+1, D^{k-1} = q\} = 1$ and collecting the terms.

The chain of equalities above shows that if $t \neq T$, then the result holds at the subsequent time period. We can use a similar argument to show that the result holds when $t = T$ as well. \blacksquare

Using Lemmas B.2 and B.3, we give a proof for Theorem 3.1.

Proof of Theorem 3.1:

Any feasible solution $(\mathbf{y}, \mathbf{u}, \mathbf{w})$ to problem (12) satisfies $y_{it}^k(q) \geq 0$, in which case, dividing both sides of the equality in Lemma B.3 by $\mathbb{P}\{D^k \geq t, D^{k-1} = q\}$, we obtain

$$\begin{aligned}
&\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \frac{\mathbb{P}\{D^k \geq t, D^{k-1} = q \mid D^\ell \geq s, D^{\ell-1} = p\}}{\mathbb{P}\{D^k \geq t, D^{k-1} = q\}} u_{js}^\ell(p) \\
&\quad - \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} \frac{\mathbb{P}\{D^k \geq t \mid D^k \geq s, D^{k-1} = q\}}{\mathbb{P}\{D^k \geq t, D^{k-1} = q\}} u_{js}^k(q) \leq c_i.
\end{aligned}$$

By the Bayes rule, the two fractions on the left side of the inequality above are, respectively, given by $\frac{\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\}}{\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p\}}$ and $\frac{\mathbb{P}\{D^k \geq s, D^{k-1} = q \mid D^k \geq t, D^{k-1} = q\}}{\mathbb{P}\{D^k \geq s, D^{k-1} = q\}}$, but for $s \leq t$, the last probability

is equal to $\frac{1}{\mathbb{P}\{D^k \geq s, D^{k-1} = q\}}$. In this case, any feasible solution to the linear program in (12) is also a feasible solution to the linear program

$$\begin{aligned}
& \max_{(\mathbf{u}, \mathbf{w}) \in \mathbb{R}_+^{|\mathcal{J}|T^2K} \times \mathcal{W}} \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{q \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j u_{jt}^k(q) & (13) \\
& \text{st} \quad \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \frac{\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\}}{\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p\}} u_{js}^\ell(p) \\
& \quad + \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} \frac{1}{\mathbb{P}\{D^k \geq s, D^{k-1} = q\}} u_{js}^k(q) \leq c_i \quad \forall i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K} \\
& \quad u_{jt}^k(q) \leq \lambda_{jt}^k w_t^k(q) \quad \forall j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K}.
\end{aligned}$$

Thus, the optimal objective value of problem (13) is an upper bound on that of problem (12), which is, in turn, an upper bound on the optimal total expected revenue.

By Lemma B.2, for any $\mathbf{w} \in \mathcal{W}$, we have $w_t^k(q) = \mathbb{P}\{D^k \geq t, D^{k-1} = q\}$. In this case, making the change of variables $x_{jt}^k(q) = \frac{1}{\mathbb{P}\{D^k \geq t, D^{k-1} = q\}} u_{jt}^k(q)$, problem (13) is equivalent to problem (2). ■

Appendix C: Fluid Approximation Through Linear Value Function Approximations

The proof for Theorem 3.1 that we give in Appendix B is based on relaxing the capacity constraints in the dynamic program in (1) through Lagrangian relaxation. In this section, we give an alternative proof for Theorem 3.1 using linear approximations to the value functions in the dynamic program. Using this approach, we also compare our fluid approximation with a linear programming approximation used by Jiang (2023). Because the remaining capacity of resource i takes values in $\{0, \dots, c_i\}$, we use $\mathcal{C} = \prod_{i \in \mathcal{L}} \{0, \dots, c_i\}$ to denote the set of all possible remaining capacity vectors. We can formulate the dynamic program in (1) as a linear program. In this linear program, we have one decision variable for each possible value of the state variable at each time period in each stage, whereas we have one constraint for each possible action that we can take at each possible value of the state variable at each time period in each state. Using the decision variables $\mathbf{J} = (J_t^k(\mathbf{y}, q) : \mathbf{y} \in \mathcal{C}, t, q \in \mathcal{T}, k \in \mathcal{K})$, as well as recalling that we use OPT to denote the optimal total expected revenue, the linear program corresponding to (1) is given by

$$\begin{aligned}
\text{OPT} &= \min_{\mathbf{J} \in \mathbb{R}^{T^2K|\mathcal{C}|}} J_1^1(\mathbf{c}, \widehat{D}^0) & (14) \\
& \text{st} \quad J_t^k(\mathbf{y}, q) \geq \sum_{j \in \mathcal{J}} \lambda_{jt}^k \left\{ f_j u_j + \theta_t^k(q) J_{t+1}^k(\mathbf{y} - \mathbf{a}_j u_j, q) + (1 - \theta_t^k(q)) J_1^{k+1}(\mathbf{y} - \mathbf{a}_j u_j, t) \right\} \\
& \quad \forall \mathbf{y} \in \mathcal{C}, \mathbf{u} \in \mathcal{F}(\mathbf{y}), t, q \in \mathcal{T}, k \in \mathcal{K},
\end{aligned}$$

where we follow the convention that $J_1^{K+1}(\mathbf{y}, q)$ is fixed at zero for all $\mathbf{y} \in \mathcal{C}$ and $q \in \mathcal{T}$. We use a linear value function approximation of the form $\widetilde{J}_t^k(\mathbf{y}, q) = \beta_t^k(q) + \sum_{i \in \mathcal{L}} \alpha_{it}^k(q) y_i$. To choose the

slope and intercept parameters $\boldsymbol{\alpha} = (\alpha_{it}^k(q) : i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K})$ and $\boldsymbol{\beta} = (\beta_t^k(q) : t, q \in \mathcal{T}, k \in \mathcal{K})$, we plug the linear value function approximation into (14) to obtain the linear program

$$\begin{aligned} \bar{Z}_{\text{linear}} = & \min_{(\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{R}^{T^2 K(|\mathcal{L}|+1)}} \beta_1^1(\hat{D}^0) + \sum_{i \in \mathcal{L}} \alpha_{i1}^1(\hat{D}^0) c_i & (15) \\ \text{st } & \beta_t^k(q) + \sum_{i \in \mathcal{L}} \alpha_{it}^k(q) y_i \\ & \geq \sum_{j \in \mathcal{J}} \lambda_{jt}^k u_j \left\{ f_j - \theta_t^k(q) \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q) a_{ij} - (1 - \theta_t^k(q)) \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t) a_{ij} \right\} \\ & \quad + \theta_t^k(q) \left\{ \beta_{t+1}^k(q) + \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q) y_i \right\} + (1 - \theta_t^k(q)) \left\{ \beta_1^{k+1}(t) + \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t) y_i \right\} \\ & \quad \forall \mathbf{y} \in \mathcal{C}, \mathbf{u} \in \mathcal{F}(\mathbf{y}), t, q \in \mathcal{T}, k \in \mathcal{K}, \end{aligned}$$

where we follow the convention that $\beta_1^{K+1}(q)$ and $\alpha_{i1}^{K+1}(q)$ are fixed at zero for all $q \in \mathcal{T}$ and $i \in \mathcal{L}$. In the constraint above, we use the fact that $\tilde{J}_t^k(\mathbf{y}, q)$ is linear in \mathbf{y} , in which case, we have $\tilde{J}_{t+1}^k(\mathbf{y} - \mathbf{a}_j u_j, q) = \beta_{t+1}^k(q) + \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q) y_i - \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q) a_{ij} u_j$ and $\tilde{J}_1^{k+1}(\mathbf{y} - \mathbf{a}_j u_j, t) = \beta_1^{k+1}(t) + \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t) y_i - \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t) a_{ij} u_j$, as well as the fact that $\sum_{j \in \mathcal{J}} \lambda_{jt}^k = 1$. We obtain the linear program in (15) by constraining the value functions in the linear program in (14) to be linear in the remaining capacities. Therefore, noting that we minimize the objective function in the linear program in (14), the optimal objective value of problem (15) cannot be smaller than the optimal objective value of problem (14), so we have $\bar{Z}_{\text{linear}} \geq \text{OPT}$. The number of constraints in (15) is exponential in the numbers of resources and products. To obtain our fluid approximation, we use a relaxation of this linear program with the number of constraints polynomial in the numbers of resources and products. In particular, arranging the terms in the constraint in (15), we write the linear program in (15) equivalently as

$$\begin{aligned} \bar{Z}_{\text{linear}} = & \min_{(\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{R}^{T^2 K(|\mathcal{L}|+1)}} \beta_1^1(\hat{D}^0) + \sum_{i \in \mathcal{L}} \alpha_{i1}^1(\hat{D}^0) c_i & (16) \\ \text{st } & \beta_t^k(q) - \theta_t^k(q) \beta_{t+1}^k(q) - (1 - \theta_t^k(q)) \beta_1^{k+1}(t) \\ & \geq \max_{\mathbf{y} \in \mathcal{C}, \mathbf{u} \in \mathcal{F}(\mathbf{y})} \left\{ \sum_{j \in \mathcal{J}} \lambda_{jt}^k u_j \left\{ f_j - \theta_t^k(q) \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q) a_{ij} - (1 - \theta_t^k(q)) \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t) a_{ij} \right\} \right. \\ & \quad \left. + \sum_{i \in \mathcal{L}} y_i \left\{ \theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t) - \alpha_{it}^k(q) \right\} \right\} \quad \forall t, q \in \mathcal{T}, k \in \mathcal{K}, \end{aligned}$$

We relax the sets \mathcal{C} and $\mathcal{F}(\mathbf{y})$, in which case, the optimal objective value of the maximization problem in the constraint becomes larger, so the constraint becomes tighter.

Letting $\tilde{\mathcal{C}} = \prod_{i \in \mathcal{L}} [0, c_i]$, we have $\tilde{\mathcal{C}} \supseteq \mathcal{C}$. Noting that $\mathcal{F}(\mathbf{y}) = \{\mathbf{u} \in \{0, 1\}^{|\mathcal{J}|} : a_{ij} u_j \leq y_i \forall i \in \mathcal{L}, j \in \mathcal{J}\}$, multiplying the constraint corresponding to product j with λ_{jt}^k and adding over all products, we

define $\tilde{\mathcal{F}}_t^k(\mathbf{y}) = \{\mathbf{u} \in \{0, 1\}^{|\mathcal{J}|} : \sum_{j \in \mathcal{J}} \lambda_{jt}^k a_{ij} u_j \leq y_i \quad \forall i \in \mathcal{L}\}$, in which case, $\tilde{\mathcal{F}}_t^k(\mathbf{y}) \supseteq \mathcal{F}(\mathbf{y})$. Thus, we can upper bound on the optimal objective value of (16) by using the linear program

$$\begin{aligned} \tilde{Z}_{\text{linear}} &= \min_{(\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{R}^{T^2 K(|\mathcal{L}|+1)}} \beta_1^1(\hat{D}^0) + \sum_{i \in \mathcal{L}} \alpha_{i1}^1(\hat{D}^0) c_i & (17) \\ \text{st } &\beta_t^k(q) - \theta_t^k(q) \beta_{t+1}^k(q) - (1 - \theta_t^k(q)) \beta_1^{k+1}(t) \\ &\geq \max_{\mathbf{y} \in \tilde{\mathcal{C}}, \mathbf{u} \in \tilde{\mathcal{F}}_t^k(\mathbf{y})} \left\{ \sum_{j \in \mathcal{J}} \lambda_{jt}^k u_j \left\{ f_j - \theta_t^k(q) \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q) a_{ij} - (1 - \theta_t^k(q)) \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t) a_{ij} \right\} \right. \\ &\quad \left. + \sum_{i \in \mathcal{L}} y_i \left\{ \theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t) - \alpha_{it}^k(q) \right\} \right\} \quad \forall t, q \in \mathcal{T}, k \in \mathcal{K}. \end{aligned}$$

Because $\tilde{\mathcal{C}} \supseteq \mathcal{C}$ and $\tilde{\mathcal{F}}_t^k(\mathbf{y}) \supseteq \mathcal{F}(\mathbf{y})$, the constraint in (17) is at least as tight as the one in (16), so $\tilde{Z}_{\text{linear}} \geq \bar{Z}_{\text{linear}}$. We can show that there exists an optimal solution to problem (17), where the decision variables $(\alpha_{it}^k(q) : i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K})$ satisfy $\alpha_{it}^k(q) \geq \theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t)$ for all $i \in \mathcal{L}, t, q \in \mathcal{T}$ and $k \in \mathcal{K}$. Theorem 2 in Adelman (2007) shows this result when there is a single stage with a fixed number of customer arrivals, but we can use the same proof technique in our problem setting. In particular, letting if $(\bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}})$ is an optimal solution to (17) with $\bar{\alpha}_{it}^k(q) < \theta_t^k(q) \bar{\alpha}_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \bar{\alpha}_{i1}^{k+1}(t)$ for some $i \in \mathcal{L}, t, q \in \mathcal{T}$ and $k \in \mathcal{K}$, then we can increase $\bar{\alpha}_{it}^k(q)$ by ϵ and decrease $\bar{\beta}_t^k(q)$ by $c_i \epsilon$ for small enough $\epsilon > 0$ to obtain a feasible solution to (17) that provides an objective value that is at least as small as the optimal objective value. Thus, we can add the constraint $\alpha_{it}^k(q) \geq \theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t)$ for all $i \in \mathcal{L}, t, q \in \mathcal{T}$ and $k \in \mathcal{K}$ to problem (17) without changing its optimal objective value, in which case, the coefficient of y_i in the maximization problem in the constraint is non-positive. Thus, we need to choose the value of y_i as small as possible. By the definition of $\tilde{\mathcal{F}}_t^k(\mathbf{y})$, the smallest value of the decision variable y_i is $\sum_{j \in \mathcal{J}} \lambda_{jt}^k a_{ij} u_j$. Thus, the maximization problem in the constraint in (17) is equivalent to

$$\begin{aligned} &\max_{\mathbf{u} \in \{0, 1\}^{|\mathcal{J}|}} \left\{ \sum_{j \in \mathcal{J}} \lambda_{jt}^k u_j \left\{ f_j - \theta_t^k(q) \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q) a_{ij} - (1 - \theta_t^k(q)) \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t) a_{ij} \right\} \right. \\ &\quad \left. + \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{J}} \lambda_{jt}^k a_{ij} u_j \left\{ \theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t) - \alpha_{it}^k(q) \right\} \right\} \\ &= \max_{\mathbf{u} \in \{0, 1\}^{|\mathcal{J}|}} \left\{ \sum_{j \in \mathcal{J}} \lambda_{jt}^k u_j \left\{ f_j - \sum_{i \in \mathcal{L}} a_{ij} \alpha_{it}^k(q) \right\} \right\} = \sum_{j \in \mathcal{J}} \lambda_{jt}^k \left[f_j - \sum_{i \in \mathcal{L}} a_{ij} \alpha_{it}^k(q) \right]^+, \end{aligned}$$

where the first equality follows by arranging the terms and the second equality holds by noting the set of feasible solutions in the maximization problem on the left side of the equality.

We replace the optimal objective value of the maximization problem in the constraint in (17) with $\sum_{j \in \mathcal{J}} \lambda_{jt}^k [f_j - \sum_{i \in \mathcal{L}} a_{ij} \alpha_{it}^k(q)]^+$. To linearize the constraint, we use the decision variable

$\eta_{jt}^k(q)$ to capture the expression $[f_j - \sum_{i \in \mathcal{L}} a_{ij} \alpha_{it}^k(q)]^+$. Recalling that we can add the constraint $\alpha_{it}^k(q) \geq \theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t)$ for all $i \in \mathcal{L}$, $t, q \in \mathcal{T}$ and $k \in \mathcal{K}$, (17) is equivalent to

$$\begin{aligned} \tilde{Z}_{\text{linear}} &= \min_{(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\eta}) \in \mathbb{R}^{T^2 K(|\mathcal{L}|+1)} \times \mathbb{R}_+^{T^2 K |\mathcal{J}|}} \beta_1^1(\hat{D}^0) + \sum_{i \in \mathcal{L}} \alpha_{i1}^1(\hat{D}^0) c_i & (18) \\ \text{st } \alpha_{it}^k(q) &\geq \theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t) \quad \forall i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K} \\ \beta_t^k(q) - \theta_t^k(q) \beta_{t+1}^k(q) - (1 - \theta_t^k(q)) \beta_1^{k+1}(t) &\geq \sum_{j \in \mathcal{J}} \lambda_{jt}^k \eta_{jt}^k(q) \quad \forall t, q \in \mathcal{T}, k \in \mathcal{K} \\ \eta_{jt}^k(q) &\geq f_j - \sum_{i \in \mathcal{L}} a_{ij} \alpha_{it}^k(q) \quad \forall j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K}, \end{aligned}$$

where we use the decision variables $\boldsymbol{\eta} = (\eta_{jt}^k(q) : j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$. The objective function in (18) matches that in (11). Viewing $\mu_{it}^k(q)$ as the slack variable for the first constraint in problem (11) and expressing this constraint as an inequality constraint by dropping the slack variable, the first constraint in (18) matches that in (11). We can show that there exists an optimal solution to problem (18) such that the second constraint is satisfied as equality. In particular, if $(\bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}}, \bar{\boldsymbol{\eta}})$ is an optimal solution to (18) with $\bar{\beta}_t^k(q) - \theta_t^k(q) \bar{\beta}_{t+1}^k(q) - (1 - \theta_t^k(q)) \bar{\beta}_1^{k+1}(t) > \sum_{j \in \mathcal{J}} \lambda_{jt}^k \bar{\eta}_{jt}^k(q)$ for some $t, q \in \mathcal{T}$ and $k \in \mathcal{K}$, then we can decrease $\bar{\beta}_t^k(q)$ by ϵ for some small enough $\epsilon > 0$ to obtain a feasible solution to (18) that provides an objective value that is at least as small as the optimal objective value. Thus, the second constraint in (18) matches that in (11). Lastly, the third constraint in (18) matches that in (11). Therefore, problems (11) and (18) are equivalent to each other. Because we obtain our fluid approximation through the dual of problem (11), the upper bound from our fluid approximation is $\tilde{Z}_{\text{linear}}$ and we can obtain our fluid approximation by using linear value function approximations. It turns out the linear program in Jiang (2023) is based on a further relaxation of the maximization problem in the constraint in (17). We have $\tilde{\mathcal{F}}_t^k(\mathbf{c}) \supseteq \tilde{\mathcal{F}}_t^k(\mathbf{y})$. Also, we have

$$\begin{aligned} \max_{\mathbf{y} \in \tilde{\mathcal{C}}, \mathbf{u} \in \tilde{\mathcal{F}}_t^k(\mathbf{c})} & \left\{ \sum_{j \in \mathcal{J}} \lambda_{jt}^k u_j \left\{ f_j - \theta_t^k(q) \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q) a_{ij} - (1 - \theta_t^k(q)) \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t) a_{ij} \right\} \right. \\ & \left. + \sum_{i \in \mathcal{L}} y_i \left\{ \theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t) - \alpha_{it}^k(q) \right\} \right\} \\ &= \sum_{j \in \mathcal{J}} \lambda_{jt}^k \left[f_j - \theta_t^k(q) \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q) a_{ij} - (1 - \theta_t^k(q)) \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t) a_{ij} \right]^+ \\ & \quad + \sum_{i \in \mathcal{L}} c_i \left[\theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t) - \alpha_{it}^k(q) \right]^+, \quad (19) \end{aligned}$$

where the equality holds by $\tilde{\mathcal{F}}_t^k(\mathbf{c}) = \{0, 1\}^{|\mathcal{J}|}$, as well as by the fact that if the coefficient of y_i in the maximization problem above is non-negative, then it is optimal to set this variable to c_i .

Because the maximization problem in (19) is a relaxation of the maximization problem in the constraint in (17), replacing the maximization problem in the constraint in (17) with the optimal

objective value of the maximization problem in (19), we can obtain an upper bound on the optimal objective value of the linear program in (17) through the linear program

$$\begin{aligned}
\widehat{Z}_{\text{linear}} &= \min_{(\alpha, \beta) \in \mathbb{R}^{T^2 K(|\mathcal{L}|+1)}} \beta_1^1(\widehat{D}^0) + \sum_{i \in \mathcal{L}} \alpha_{i1}^1(\widehat{D}^0) c_i & (20) \\
\text{st } & \beta_t^k(q) - \theta_t^k(q) \beta_{t+1}^k(q) - (1 - \theta_t^k(q)) \beta_1^{k+1}(t) \\
& \geq \sum_{j \in \mathcal{J}} \lambda_{jt}^k \left[f_j - \theta_t^k(q) \sum_{i \in \mathcal{L}} \alpha_{i,t+1}^k(q) a_{ij} - (1 - \theta_t^k(q)) \sum_{i \in \mathcal{L}} \alpha_{i1}^{k+1}(t) a_{ij} \right]^+ \\
& \quad + \sum_{i \in \mathcal{L}} c_i \left[\theta_t^k(q) \alpha_{i,t+1}^k(q) + (1 - \theta_t^k(q)) \alpha_{i1}^{k+1}(t) - \alpha_{it}^k(q) \right]^+ \quad \forall t, q \in \mathcal{T}, k \in \mathcal{K}.
\end{aligned}$$

Because the constraint in (20) is at least as tight as the constraint in (17), the optimal objective values of these two linear programs satisfy $\widehat{Z}_{\text{linear}} \geq \widetilde{Z}_{\text{linear}}$. Using the fact that $\widetilde{Z}_{\text{linear}} \geq \overline{Z}_{\text{linear}} \geq \text{OPT}$, we get $\widehat{Z}_{\text{linear}} \geq \widetilde{Z}_{\text{linear}} \geq \overline{Z}_{\text{linear}} \geq \text{OPT}$. Thus, noting that the optimal objective value of problem (18) is $\widetilde{Z}_{\text{linear}}$, the linear programs in (18) and (20) both provide upper bounds on the optimal total expected revenue, but the upper bound provided by (18) is at least as tight as the one provided by (20). The linear program in (20) is precisely the analogue of the linear program given in Section 3 of Jiang (2023) when the demands occur in multiple stages, there is a random number of customer arrivals in each stage and the numbers of customer arrivals in successive stages are dependent on each other. The author uses this linear program to compute an upper bound on the optimal total expected revenue, but does not construct an approximate policy by using an optimal solution to this linear program. By the preceding discussion, the upper bound provided by our fluid approximation is at least as tight as the one provided by the linear program in Jiang (2023).

Appendix D: Verifying the Fluid Approximation Through Decisions of the Optimal Policy

We use the decisions made by the optimal policy to construct a feasible solution to the problem (2) in such a way that the objective value provided by this solution for problem (2) matches the optimal total expected revenue. We use the Bernoulli random variable X_{jt}^k to capture the decision of the optimal policy for product j at time period t in stage k , where $X_{jt}^k = 1$ if and only if the optimal policy accepts a request for product j at time period t in stage k . By the dynamic program in (1), given that we reach time period t in stage k before this stage is over, the random variable X_{jt}^k depends on the remaining capacities of the resources at time period t in stage k and the demand in stage $k-1$. Given that we reach time period t in stage k , the remaining capacities of the resources at time period t in stage k depend on the product requests at all time periods up to time period t in stage k , as well as the demand random variables in stages up to stage k . Thus, given that we reach time period t in stage k before this stage is over, the random variable X_{jt}^k depends on the demand random variables D^1, \dots, D^{k-1} , but not on the other demand random variables. Given

that we do not reach time period t in stage k , we have $X_{jt}^k = 0$. By the preceding discussion, the random variable X_{jt}^k depends on the demand random variables D^1, \dots, D^{k-1} , as well as $\mathbf{1}(D^k \geq t)$, but not on the other demand random variables. Letting $\bar{x}_{jt}^k(q) = \mathbb{E}\{X_{jt}^k | D^k \geq t, D^{k-1} = q\}$, we show that the solution $\bar{\mathbf{x}} = (\bar{x}_{jt}^k(q) : j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$ provides an objective value for problem (2) that is equal to the optimal total expected revenue. The optimal total expected revenue is $\text{OPT} = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j \mathbb{E}\{X_{jt}^k\}$. Given $D^k < t$ so that we do not reach time period t in stage k , we have $X_{jt}^k = 0$, so by the tower property of conditional expectations, we obtain $\mathbb{E}\{X_{jt}^k\} = \sum_{q \in \mathcal{T}} \mathbb{E}\{X_{jt}^k | D^k \geq t, D^{k-1} = q\} \mathbb{P}\{D^k \geq t, D^{k-1} = q\}$. By the definition of $\bar{x}_{jt}^k(q)$, the last equality is equivalent to $\mathbb{E}\{X_{jt}^k\} = \sum_{q \in \mathcal{T}} \bar{x}_{jt}^k(q) \mathbb{P}\{D^k \geq t, D^{k-1} = q\}$. Therefore, we can express the optimal total expected revenue as $\text{OPT} = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j \sum_{q \in \mathcal{T}} \bar{x}_{jt}^k(q) \mathbb{P}\{D^k \geq t, D^{k-1} = q\}$, in which case, noting the objective function of problem (2), the objective value provided by the solution $\bar{\mathbf{x}}$ for problem (2) is equal to the optimal total expected revenue.

We show that the solution $\bar{\mathbf{x}}$ is feasible to problem (2). We use the Bernoulli random variable Λ_{jt}^k to capture whether there is a request for product j at time period t in stage k , where $\Lambda_{jt}^k = 1$ if and only if there is a request for product j at time period t in stage k . The Bernoulli random variable Λ_{jt}^k has parameter λ_{jt}^k and it is independent of the demand random variables. If $D^k < t$, then we have $X_{jt}^k = 0$, which implies that $X_{jt}^k \leq \mathbf{1}(D^k \geq t)$. On the other hand, we can accept a request for a product only if there is request for it, which implies that $X_{jt}^k \leq \Lambda_{jt}^k$. In this case, we obtain $X_{jt}^k \leq \mathbf{1}(D^k \geq t, \Lambda_{jt}^k = 1)$. If we take the expectation of both sides of the last inequality conditional on $D^k \geq t$ and $D^{k-1} = q$, then we obtain $\bar{x}_{jt}^k(q) \leq \mathbb{E}\{\Lambda_{jt}^k | D^k \geq t, D^{k-1} = q\} = \lambda_{jt}^k$, where the equality holds because the random variable Λ_{jt}^k is independent of D^{k-1} and D^k . Thus, the solution $\bar{\mathbf{x}}$ satisfies the second constraint in (2). We give a useful identity. Given that $D^\ell < s$, so that we do not reach time period s in stage ℓ , we have $X_{js}^\ell = 0$. Therefore, considering some stage $\ell = 1, \dots, k-1$, by the tower property of conditional expectations, we get

$$\begin{aligned} & \mathbb{E}\{X_{js}^\ell | D^k \geq t, D^{k-1} = q\} \\ &= \sum_{p \in \mathcal{T}} \mathbb{E}\{X_{js}^\ell | D^\ell \geq s, D^{\ell-1} = p, D^k \geq t, D^{k-1} = q\} \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p | D^k \geq t, D^{k-1} = q\} \\ &\stackrel{(a)}{=} \sum_{p \in \mathcal{T}} \mathbb{E}\{X_{js}^\ell | D^\ell \geq s, D^{\ell-1} = p\} \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p | D^k \geq t, D^{k-1} = q\}, \\ &= \sum_{p \in \mathcal{T}} \bar{x}_{js}^\ell(p) \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p | D^k \geq t, D^{k-1} = q\}, \end{aligned} \tag{21}$$

where (a) holds because X_{js}^ℓ depends on the demand random variables $D^1, \dots, D^{\ell-1}$, as well as $\mathbf{1}(D^\ell \geq s)$, but given that $D^{\ell-1} = p$, $D^1, \dots, D^{\ell-1}$ are independent of D^{k-1} and D^k .

The total capacity consumption of resource i up to time period t in stage k cannot exceed the capacity of the resource, so we have $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} X_{js}^\ell + \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} X_{js}^k \leq c_i$ with

probability one. Taking the expectation of both sides of the last inequality conditional on $D^k \geq t$ and $D^{k-1} = q$, we obtain the chain of inequalities

$$\begin{aligned} c_i &\geq \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{E}\{X_{js}^\ell \mid D^k \geq t, D^{k-1} = q\} + \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} \mathbb{E}\{X_{js}^k \mid D^k \geq t, D^{k-1} = q\} \\ &\stackrel{(b)}{=} \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \sum_{p \in \mathcal{T}} \bar{x}_{js}^\ell(p) \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\} + \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} \bar{x}_{js}^k(q), \end{aligned}$$

where (b) uses (21) and the fact that $\mathbb{E}\{X_{js}^k \mid D^k \geq t, D^{k-1} = q\} = \mathbb{E}\{X_{js}^k \mid D^k \geq s, D^{k-1} = q\} = \bar{x}_{js}^k(q)$ because X_{js}^k depends on $\mathbf{1}(D^k \geq s), D^{k-1}, \dots, D^1$. Thus, the solution $\bar{\mathbf{x}}$ satisfies the first constraint in (2) as well. By the discussion in this section, the solution $\bar{\mathbf{x}}$ is feasible to problem (2) and provides an objective value for this problem that is equal to the optimal total expected revenue. In this case, the optimal objective value of problem (2) is at least as large as the optimal total expected revenue. This discussion gives an alternative proof of the fact that the optimal objective value of our fluid approximation is an upper bound on the optimal total expected revenue. This discussion is not constructive in the sense that it does not allow us to derive the form of problem (2). Instead, it verifies that the optimal objective value of problem (2) is an upper bound on the optimal total expected revenue once we are given the form of problem (2). On the other hand, using the Lagrangian relaxation strategy as in Appendix B or using the linear value function approximations as in Appendix C allows us to derive the form of problem (2), while also establishing that problem (2) provides an upper bound on the optimal total expected revenue.

Appendix E: Comparing the Fluid Approximation and the Offline Bound

We give problem instances to demonstrate that the upper bound on the optimal total expected revenue provided by our fluid approximation can be tighter or looser than the offline bound. First, we give a problem instance for which the upper bound provided by our fluid approximation is tighter than the offline bound. There are two stages. The support of the demand in each stage is one. There is one resource and two products. Thus, we have $K = 1$, $T = 1$, $\mathcal{L} = \{1\}$ and $\mathcal{J} = \{1, 2\}$. The capacity of the resource is $c_1 = 1$. The revenues of the products are $f_1 = 1$ and $f_2 = 2$. Fixing the demand in the stage right before the beginning of the selling horizon at \widehat{D}^0 , the distribution of the demand random variables in each stage is characterized by the conditional probabilities $\mathbb{P}\{D^1 = 0 \mid D^0 = \widehat{D}^0\} = \mathbb{P}\{D^1 = 1 \mid D^0 = \widehat{D}^0\} = \frac{1}{2}$, $\mathbb{P}\{D^2 = 1 \mid D^1 = 0\} = 1$ and $\mathbb{P}\{D^2 = 0 \mid D^1 = 1\} = \mathbb{P}\{D^2 = 1 \mid D^1 = 1\} = \frac{1}{2}$. Lastly, the probability of getting a request for each product at each time period in each stage is given by $\lambda_{11}^1 = 1$, $\lambda_{11}^2 = \frac{1}{2}$ and $\lambda_{21}^2 = \frac{1}{2}$. Considering the offline bound, because the capacity of the resource is one and the revenue of the second product is larger than that of the first product, if there is a request for the second product at any time period in any stage, then we

accept the request for the second product. If there is a request for the first product but not the second product, then we accept the request for the first product. The probability that we have a request for the second product in the first stage is zero. To have a request for the second product in the second stage, we must have a demand of one in the second stage and this demand must be for the second product. Thus, the probability that we have a request for the second product is $(\frac{1}{2} + \frac{1}{2} \times \frac{1}{2}) \times \frac{1}{2} = \frac{3}{8}$. On the other hand, the probability that we have a request for the first product but not the second product is given by $\frac{1}{2} \times \frac{1}{2} + \frac{1}{2} \times \frac{1}{2} + \frac{1}{2} \times \frac{1}{2} \times \frac{1}{2} = \frac{5}{8}$, where the three terms in the sum, respectively, correspond to having one demand in the first stage and no demand in the second stage, no demand in the first stage and one demand in the second stage and one demand in both stages, along with a demand for the first product and no demand for the second product. Thus, the offline bound is $\bar{Z}_{\text{offline}} = \frac{3}{8} \times 2 + \frac{5}{8} \times 1 = \frac{11}{8}$. Considering problem (2), given that the demand in the previous stage is q , if we never have a request for product j at time period t in stage k , then we can drop the decision variable $x_{jt}^k(q)$. For our problem instance, problem (2) is given by

$$\begin{aligned} \bar{Z}_{\text{LP}} = \max \quad & \frac{1}{2} x_{11}^1(\widehat{D}^0) + \frac{1}{2} x_{11}^2(0) + \frac{1}{2} 2 x_{21}^2(0) + \frac{1}{4} x_{11}^1(1) + \frac{1}{4} 2 x_{21}^2(1) \\ \text{st} \quad & x_{11}^1(\widehat{D}^0) + x_{11}^2(1) + x_{21}^2(1) \leq 1 \\ & x_{11}^2(0) + x_{21}^2(0) \leq 1 \\ & x_{11}^1(\widehat{D}^0) \in [0, 1], x_{11}^2(0), x_{21}^2(0), x_{11}^1(1), x_{21}^2(1) \in [0, \frac{1}{2}]. \end{aligned}$$

The optimal objective value of the problem above is $\frac{5}{4}$. Therefore, we have $\bar{Z}_{\text{LP}} = \frac{5}{4} \leq \frac{11}{8} = \bar{Z}_{\text{offline}}$, so the upper bound from our fluid approximation can be tighter than the offline bound.

Second, we give a problem instance for which the upper bound provided by our fluid approximation is looser than the offline bound. This problem instance has the same parameters as the problem instance in the previous paragraph other than the distribution of the demand random variables in different stages and the request probabilities for different products at different time periods in different stages. In particular, we have $\mathbb{P}\{D^1 = 1 \mid D^0 = \widehat{D}^0\} = 1$, $\mathbb{P}\{D^2 = 1 \mid D^1 = 1\} = 1$, so the demand in each stage is always one. Furthermore, we have $\lambda_{11}^1 = \lambda_{21}^1 = \frac{1}{2}$ and $\lambda_{21}^2 = \frac{1}{2}$, so we have a request for each of the two products at the first time period in the first stage with equal probabilities, whereas we have a request for the second product at the first time period in the second stage with probability $\frac{1}{2}$ and we do not have a request at the first time period in the second stage with probability $\frac{1}{2}$. To ensure that we have a product request at each time period in each stage with probability one, we can introduce a third product with a revenue of zero such that we have a request for the product at the first time period in the second stage with probability $\frac{1}{2}$. To keep our problem instance succinct, we do not explicitly introduce such a product. We have a request for

the second product at some period in some stage with probability $\frac{1}{2} \times \frac{1}{2} + \frac{1}{2} \times \frac{1}{2} + \frac{1}{2} \times \frac{1}{2} = \frac{3}{4}$, where the three terms in the sum, respectively, correspond to having a request for the second product in the first stage but not in the second stage, not in the first stage but in the second stage and in both stages. On the other hand, the probability that we have a request for the first product but not the second product is $\frac{1}{2} \times \frac{1}{2} = \frac{1}{4}$. Thus, the offline bound is $\bar{Z}_{\text{offline}} = \frac{3}{4} \times 2 + \frac{1}{4} = \frac{7}{4}$. Considering problem (2), for our problem instance, this problem is given by

$$\begin{aligned} \bar{Z}_{\text{LP}} = \max \quad & x_{11}^1(\hat{D}^0) + 2x_{21}^1(\hat{D}^0) + 2x_{21}^2(1) \\ \text{st} \quad & x_{11}^1(\hat{D}^0) + x_{21}^1(\hat{D}^0) + x_{21}^2(1) \leq 1 \\ & x_{11}^1(\hat{D}^0), x_{21}^1(\hat{D}^0), x_{21}^2(1) \in [0, \frac{1}{2}]. \end{aligned}$$

The optimal objective value of the problem above is 2. Therefore, we have $\bar{Z}_{\text{LP}} = 2 \geq \frac{7}{4} = \bar{Z}_{\text{offline}}$, so the upper bound from our fluid approximation can be looser than the offline bound.

Appendix F: Detailed Description of the Approximate Policy

In Table 1, we give a detailed description of our approximate policy. We use the variable y_{it}^k to keep track of the remaining capacity of resource i at the beginning of time period t in stage k . In Line 1, we initialize the remaining capacities of the resources, as well as set the demand in the stage right before the beginning of the selling horizon to \hat{D}^0 , which is a part of the problem data. In Line 2, the for loop enumerates over the stages in the selling horizon. We use the variable k to keep track of the current stage. In Line 3, we use the variable t to keep track of the current time period and S^k to keep track of whether the current stage should end. We initialize the current time period to one and ensure that the current stage should not end at least until after the first time period. In Line 4, the while loop enumerates over the time periods in the current stage.

In Line 5, we sample the product requested at the current time period. We use the variable J_t^k to keep track of the product requested at the current time period. In Line 6, we sample whether the approximate policy would accept the product request at the current time period. The binary variable A_t^k takes value one if and only if the approximate policy would accept the product request. In Line 7, we check whether we have enough capacity to serve the product requested at the current time period. The binary variable C_t^k takes value one if and only if we have enough capacity. In Lines 8 to 10, we check whether the approximate policy would accept the product requested at the current time period and we have enough capacity to serve the requested product. If so, then we decrease the capacities of the resources used by the requested product.

In Line 11, recalling that $\theta_t^k(q) = \mathbb{P}\{D^k \geq t+1 \mid D^k \geq t, D^k = q\}$, we sample whether there is one more time period in the current stage. The binary variable E^k takes value one if and only if there

-
- [1] Set $y_{i1}^1 = c_i$ for all $i \in \mathcal{L}$ and $D^0 = \widehat{D}^0$.
 - [2] For $k = 1, \dots, K$, do the following steps.
 - [3] Set $t = 1$, $S^k = \text{False}$.
 - [4] While $S^k = \text{False}$ do the following steps.
 - [5] Sample $J_t^k \in \mathcal{J}$ such that $\mathbb{P}\{J_t^k = j\} = \lambda_{j,t}^k$.
 - [6] Sample $A_t^k \in \{0, 1\}$ such that $\mathbb{P}\{A_t^k = 1\} = \gamma \frac{\bar{x}_{J_t^k, t}^{J_t^k, t}(D^{k-1})}{\lambda_{J_t^k, t}^k}$.
 - [7] Set $C_t^k = \mathbf{1}(y_{it}^k \geq a_{i, J_t^k} \forall i \in \mathcal{L})$.
 - [8] If $A_t^k = 1$ and $C_t^k = 1$, then
 - [9] Set $y_{i, t+1}^k = y_{it}^k - a_{i, J_t^k}$ for all $i \in \mathcal{L}$
 - [10] End if.
 - [11] Sample $E^k \in \{0, 1\}$ such that $\mathbb{P}\{E^k = 1\} = \theta_t^k(D^{k-1})$.
 - [12] If $E^k = 1$, then
 - [13] Increment t by one
 - [14] Else
 - [15] Set $S^k = \text{True}$, $D^k = t$ and $y_{i1}^{k+1} = y_{i, t+1}^k$ for all $i \in \mathcal{L}$
 - [16] End if.
 - [17] End While.
 - [18] End For.
-

Table 1 Detailed description of the approximate policy from the fluid approximation.

is one more time period in the current stage. In Lines 12 to 16, we check whether there is one more time period. If so, then we increment the current time period. Otherwise, we set the variable S^k to indicate that the current stage is to end and set the value of the demand in the current stage, as well as the remaining capacities of the resources at the beginning of the next stage.

Appendix G: Upper Bound on the Optimality Gap of the Approximate Policy

We give a problem instance such that if we scale the number of stages and resource capacities with the same rate θ , then the ratio between the optimal total expected revenue and the optimal objective value of (2) is at most $O\left(1 - \frac{1}{\sqrt{\theta}}\right)$. There are K stages with $K \geq 72\pi$. The demand in each stage is always one. There is one resource and one product. The capacity of the resource is $\frac{1}{2}K$. The revenue of the product is one. At each time period, we have a demand for the product with probability $\frac{1}{2}$. Because there is one product, it is optimal to accept the product requests as much as the capacity of the resource allows. Thus, letting B^1, \dots, B^K be independent Bernoulli random variables with parameter $\frac{1}{2}$, the optimal total expected revenue is $\text{OPT} = \mathbb{E}\{\min\{\sum_{k=1}^K B^k, \frac{1}{2}K\}\} = \mathbb{E}\{\min\{\sum_{k=1}^K (B^k - \frac{1}{2}), 0\}\} + \frac{1}{2}K$. The function $h(x) = \min\{x, 0\}$ is Lipschitz continuous with

Lipschitz constant one. We have $\sum_{k=1}^K \text{Var}(B^k - \frac{1}{2}) = \frac{1}{4}K$ and $\mathbb{E}\{|B^k - \frac{1}{2}|^3\} = \frac{1}{8}$. In this case, letting N be the standard normal random variable, by Theorem 3.1 in Chen et al. (2010), we obtain

$$\frac{\mathbb{E}\{\min\{\sum_{k=1}^K (B^k - \frac{1}{2}), 0\}\}}{\sqrt{\frac{1}{4}K}} \leq \mathbb{E}\{\min\{N, 0\}\} + \frac{3}{\sqrt{K}}.$$

Arranging the terms in the inequality above and using the identity $\mathbb{E}\{\min\{N, 0\}\} = -\frac{1}{\sqrt{2\pi}}$, we have $\mathbb{E}\{\min\{\sum_{k=1}^K (B^k - \frac{1}{2}), 0\}\} \leq -\frac{1}{2}\sqrt{\frac{K}{2\pi}} + \frac{3}{2}$. Because $\text{OPT} = \mathbb{E}\{\min\{\sum_{k=1}^K (B^k - \frac{1}{2}), 0\}\} + \frac{1}{2}K$, we get $\text{OPT} \leq \frac{1}{2}K - \frac{1}{2}\sqrt{\frac{K}{2\pi}} + \frac{3}{2}$. Noting that the demand in each stage is always one and there is one product, using the decision variables $\mathbf{x} = (x_{11}^k(1) : k \in \mathcal{K})$, problem (2) for our problem instance is given by $\bar{Z}_{\text{LP}} = \max_{\mathbf{x} \in \mathbb{R}_+^K} \{\sum_{k \in \mathcal{K}} x_{11}^k(1) : \sum_{k \in \mathcal{K}} x_{11}^k(1) \leq \frac{1}{2}K, x_{11}^k(1) \leq \frac{1}{2} \forall k \in \mathcal{K}\}$. The optimal objective value of this problem is $\bar{Z}_{\text{LP}} = \frac{1}{2}K$. Therefore, we obtain $\frac{\text{OPT}}{\bar{Z}_{\text{LP}}} \leq 1 - \frac{1}{\sqrt{2\pi K}} + \frac{3}{K} = 1 - \frac{1}{\sqrt{2\pi K}} + \frac{3}{\sqrt{K}\sqrt{K}} \leq 1 - \frac{1}{2\sqrt{2\pi K}} \leq 1 - \frac{1}{\sqrt{26K}}$, where the second inequality uses the fact that $K \geq 72\pi$. In our problem instance, the number of stages is K and the capacity of the resource is $\frac{1}{2}K$. Thus, if we scale the number of stages and capacity of the resource with rate θ , then the optimal total expected revenue and the optimal objective value of (2) is separated by $1 - \frac{1}{\sqrt{26\theta}}$.

Appendix H: Problem Instance with Increasing Support for Demands

When there is a single stage, so the demand model is not aware of the calendar and the dependence between the demands in different stages is not an issue, we can construct fluid approximations such that the relative gap between the optimal objective value of the fluid approximation and the optimal total expected revenue vanishes as the capacities of the resources get large, irrespective of how the demand is scaled; see, for example, Bai et al. (2023). We give a counterexample to demonstrate that if we scale the support of the demand along with the capacities of the resources, then the relative gap between the optimal objective value of problem (2) and the optimal total expected revenue does not necessarily vanish, even when the demands at different stages are independent of each other. We consider a problem instance with $K = 3$ stages. There is one resource with a capacity of $C + 1$ for an even integer C . There are two products indexed by $\{1, 2\}$. The revenue associated with the two products are $f_1 = 1$ and $f_2 = C/4$. The demand in the first stage can take two values with $\mathbb{P}\{D^1 = 0\} = 1/2$ and $\mathbb{P}\{D^1 = C\} = 1/2$. The demand in the second and third stages have the distributions $\mathbb{P}\{D^2 = 1\} = 1$ and $\mathbb{P}\{D^3 = C/2\} = 1$. The probabilities of getting requests for the different products are given by $\lambda_{1t}^1 = 1$ for all $t = 1, \dots, C$, $\lambda_{11}^2 = 1$, $\lambda_{1t}^3 = 1$ for all $t = 1, \dots, C/2 - 1$ and $\lambda_{2, C/2}^3 = 1$. All other request probabilities are zero. Thus, we have a request for the second product only at the last time period in the last stage.

The optimal policy for this problem instance is to accept all product requests until there is one unit of remaining capacity and save that unit of capacity for the request for the second product. In

this case, if the demand in the first stage is C , then we obtain a total revenue of $C + \frac{1}{4}C = \frac{5}{4}C$. If the demand in the first stage is zero, then we obtain a total revenue of $\frac{1}{2}C + \frac{1}{4}C = \frac{3}{4}C$. Therefore, the optimal total expected revenue is C , so $\text{OPT} = C$. In the linear program in (2), noting the last constraint, if $\lambda_{jt}^k = 0$, then we can drop the decision variables $\{x_{jt}^k(q) : q \in \mathcal{T}\}$. In this case, using \widehat{D}^0 to denote the demand right before the beginning of the selling horizon, for the problem instance in the previous paragraph, the linear program in (2) is given by

$$\begin{aligned} \overline{Z}_{\text{LP}} = \max \quad & \frac{1}{2} \sum_{t=1}^C x_{1t}^1(\widehat{D}^0) + \frac{1}{2} x_{11}^2(0) + \frac{1}{2} x_{11}^2(C) + \sum_{t=1}^{C/2-1} x_{1t}^3(1) + \frac{1}{4} C x_{2,C/2}^3(1) \\ \text{st} \quad & \sum_{t=1}^C x_{1t}^1(\widehat{D}^0) \leq C + 1 \\ & x_{11}^2(0) \leq C + 1 \\ & \sum_{t=1}^C x_{1t}^1(\widehat{D}^0) + x_{11}^2(C) \leq C + 1 \\ & \frac{1}{2} \sum_{t=1}^C x_{1t}^1(\widehat{D}^0) + \frac{1}{2} x_{11}^2(0) + \frac{1}{2} x_{11}^2(C) + \sum_{t=1}^{C/2-1} x_{1t}^3(1) + x_{2,C/2}^3(1) \leq C + 1, \end{aligned}$$

where we understand that all decision variables are restricted to be in $[0, 1]$. Setting all decision variables to one yields a feasible solution, so $\overline{Z}_{\text{LP}} = \frac{1}{2}C + 1 + \frac{1}{2}C - 1 + \frac{1}{4}C = \frac{5}{4}C$.

Thus, for this problem instance, we have $\frac{\text{OPT}}{\overline{Z}_{\text{LP}}} = \frac{4}{5}$. Therefore, the ratio $\frac{\text{OPT}}{\overline{Z}_{\text{LP}}}$ stays away from one for this problem instance as the resource capacity gets large.

Appendix I: Performance Guarantee for the Approximate Policy

We show that the total expected revenue of the approximate policy satisfies $\frac{\text{APX}}{\text{OPT}} \geq \frac{\text{APX}}{\overline{Z}_{\text{LP}}} \geq \frac{1}{4L}$. We use the random variables $\Psi_t^k(q)$, G_{jt}^k and $N_{it}^k(q)$ as defined at the beginning of Section 5. By the discussion right after (3), if we can show that $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} \geq \alpha$, then $\text{APX} \geq \gamma \alpha \overline{Z}_{\text{LP}}$. By (4), to lower bound the probability $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\}$, it is enough to upper bound the probability $\mathbb{P}\{\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p) + \sum_{s=1}^t N_{is}^k(q) \geq c_i \mid D^{k-1} = q\}$. Recall that the random variables $\{N_{it}^k(q) : i \in \mathcal{L}, t \in \mathcal{T}, k \in \mathcal{K}\}$ are independent of demands, so we get

$$\begin{aligned} & \mathbb{E} \left\{ \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p) + \sum_{s=1}^t N_{is}^k(q) \mid D^{k-1} = q \right\} \\ & \stackrel{(a)}{=} \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^{k-1} = q\} \sum_{j \in \mathcal{J}} a_{ij} \gamma \bar{x}_{js}^\ell(p) + \sum_{s=1}^t \sum_{j \in \mathcal{J}} a_{ij} \gamma \bar{x}_{js}^k(q) \stackrel{(b)}{\leq} \gamma c_i, \end{aligned}$$

where (a) holds by using the definition of $\Psi_t^k(p)$ and $\mathbb{E}\{N_{it}^k(q)\} = \sum_{j \in \mathcal{J}} a_{ij} \bar{x}_{jt}^k(q)$, whereas (b) holds by noting that \bar{x} satisfies the first constraint in problem (2), as well as using the fact that

conditional on D^{k-1} , D^k is independent of D^1, \dots, D^{k-1} , in which case, for $\ell \leq k-1$, we have $\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^{k-1} = q\} = \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\}$. Therefore, we get

$$\begin{aligned} & \mathbb{P}\left\{\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p) + \sum_{s=1}^t N_{is}^k(q) \geq c_i \mid D^{k-1} = q\right\} \\ & \stackrel{(c)}{\leq} \frac{1}{c_i} \mathbb{E}\left\{\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^\ell(p) N_{is}^\ell(p) + \sum_{s=1}^t N_{is}^k(q) \geq c_i \mid D^{k-1} = q\right\} \leq \gamma, \end{aligned}$$

where (c) is the Markov inequality. Thus, by (4), we have $\mathbb{P}\{G_{jt}^k = 1 \mid D^{k-1} = q\} \geq 1 - L\gamma$, yielding $\text{APX} \geq \gamma(1 - L\gamma)\bar{Z}_{\text{LP}}$. Setting $\gamma = \frac{1}{2L}$ and using Theorem 3.1, we get $\frac{1}{4L} \leq \frac{\text{APX}}{\bar{Z}_{\text{LP}}} \leq \frac{\text{APX}}{\text{OPT}}$. ■

Appendix J: Auxiliary Results for Concentration Inequalities

We give proofs for two results used in Section 5. First, we show that $|M_i^k(\ell) - M_i^k(\ell+1)| \leq \frac{2}{\epsilon^3}$ with probability one. Letting $V^k(q) = \mathbb{P}\{D^k = q\}$ and $\theta = 1 - \epsilon$ for notational brevity, we define

$$Q^k(p, q) = \frac{1}{\theta} [\mathbb{P}\{D^k = q \mid D^{k-1} = p\} - (1 - \theta)V^k(q)]. \quad (22)$$

By the assumption that $\mathbb{P}\{D^k = q \mid D^{k-1} = p\} \geq \epsilon$, we have $Q^k(p, q) \geq \frac{1}{\theta}(\epsilon - (1 - \theta)) = 0$. Also, we have $\sum_{q \in \mathcal{T}} Q^k(p, q) = \frac{1}{\theta} [\sum_{q \in \mathcal{T}} \mathbb{P}\{D^k = q \mid D^{k-1} = p\} - (1 - \theta)\sum_{q \in \mathcal{T}} V^k(q)] = 1$. Thus, we can use $Q^k(p, q)$ to characterize the transition probabilities of a non-stationary Markov chain. Consider the non-stationary Markov chain Y^0, Y^1, Y^2, \dots over the state space \mathcal{T} characterized by the transition probabilities $\mathbb{P}\{Y^k = q \mid Y^{k-1} = p\} = Q^k(p, q)$ for all $p, q \in \mathcal{T}$ with $\mathbb{P}\{Y^0 = q\} = \mathbf{1}(D^0 = q)$. To show that $|M_i^{k-1}(\ell) - M_i^{k-1}(\ell+1)| \leq \frac{2}{\epsilon^3}$ with probability one, we will use two preliminary lemmas. In the next lemma, we slightly extend Theorem 4.9 in Levin and Peres (2017), which characterizes the mixing times of Markov chains, to non-stationary Markov chains.

Lemma J.1 *For all $p, q \in \mathcal{T}$ and $k, \ell \in \mathcal{K}$ with $\ell \geq k+1$, we have*

$$\mathbb{P}\{D^\ell = q \mid D^k = p\} - \mathbb{P}\{D^\ell = q\} = (1 - \epsilon)^{(\ell-k)} \left[\mathbb{P}\{Y^\ell = q \mid Y^k = p\} - \mathbb{P}\{D^\ell = q\} \right].$$

Proof: Letting $\theta = 1 - \epsilon$ for notational brevity, for all $p, q \in \mathcal{T}$ and $k, \ell \in \mathcal{K}$ with $\ell \geq k+1$, we claim that $\mathbb{P}\{D^\ell = q \mid D^k = p\} = (1 - \theta^{\ell-k})\mathbb{P}\{D^\ell = q\} + \theta^{\ell-k}\mathbb{P}\{Y^\ell = q \mid Y^k = p\}$. We show the claim by using induction over $\ell = k+1, \dots, K$. Consider the case $\ell = k+1$. By the definition of $Q^{k+1}(p, q)$, we have $\mathbb{P}\{Y^{k+1} = q \mid Y^k = p\} = Q^{k+1}(p, q) = \frac{1}{\theta} [\mathbb{P}\{D^{k+1} = q \mid D^k = p\} - (1 - \theta)\mathbb{P}\{D^{k+1} = q\}]$, so arranging the terms, we get $(1 - \theta)\mathbb{P}\{D^{k+1} = q\} + \theta\mathbb{P}\{Y^{k+1} = q \mid Y^k = p\} = \mathbb{P}\{D^{k+1} = q \mid D^k = p\}$, establishing the claim for $\ell = k+1$. Assuming that the claim holds for $\ell \geq k+1$, we show that the claim holds for $\ell+1 \geq k+1$ as well. Arranging the terms in the definition of $Q^{\ell+1}(s, q)$ in (22), we have $\mathbb{P}\{D^{\ell+1} = q \mid D^\ell = s\} = \theta Q^{\ell+1}(s, q) + (1 - \theta)V^{\ell+1}(q)$. In this case, noting the

identity $\mathbb{P}\{Y^{\ell+1} = q | Y^k = p\} = \sum_{s \in \mathcal{T}} \mathbb{P}\{Y^{\ell+1} = q | Y^\ell = s\} \mathbb{P}\{Y^\ell = s | Y^k = p\}$, it follows that we have the chain of equalities

$$\begin{aligned}
& \sum_{s \in \mathcal{T}} \mathbb{P}\{D^{\ell+1} = q | D^\ell = s\} \mathbb{P}\{Y^\ell = s | Y^k = p\} \\
&= \sum_{s \in \mathcal{T}} \left[\theta Q^{\ell+1}(s, q) + (1 - \theta) V^{\ell+1}(q) \right] \mathbb{P}\{Y^\ell = s | Y^k = p\} \\
&\stackrel{(a)}{=} \theta \sum_{s \in \mathcal{T}} \mathbb{P}\{Y^{\ell+1} = q | Y^\ell = s\} \mathbb{P}\{Y^\ell = s | Y^k = p\} + (1 - \theta) V^{\ell+1}(q) \\
&\stackrel{(b)}{=} \theta \mathbb{P}\{Y^{\ell+1} = q | Y^k = p\} + (1 - \theta) \mathbb{P}\{D^{\ell+1} = q\}, \tag{23}
\end{aligned}$$

where (a) holds because $Q^{\ell+1}(s, q) = \mathbb{P}\{Y^{\ell+1} = q | Y^\ell = s\}$ and $\sum_{s \in \mathcal{T}} \mathbb{P}\{Y^\ell = s | Y^k = p\} = 1$, whereas (b) uses the fact that we have $V^{\ell+1}(q) = \mathbb{P}\{D^{\ell+1} = q\}$ by its definition.

We have $\mathbb{P}\{D^\ell = s | D^k = p\} = (1 - \theta^{\ell-k}) \mathbb{P}\{D^\ell = s\} + \theta^{\ell-k} \mathbb{P}\{Y^\ell = s | Y^k = p\}$ by the induction assumption, so noting that $\mathbb{P}\{D^{\ell+1} = q\} = \sum_{s \in \mathcal{T}} \mathbb{P}\{D^{\ell+1} = q | D^\ell = s\} \mathbb{P}\{D^\ell = s\}$, we have

$$\begin{aligned}
\mathbb{P}\{D^{\ell+1} = q | D^k = p\} &= \sum_{s \in \mathcal{T}} \mathbb{P}\{D^{\ell+1} = q | D^\ell = s\} \mathbb{P}\{D^\ell = s | D^k = p\} \\
&= \sum_{s \in \mathcal{T}} \mathbb{P}\{D^{\ell+1} = q | D^\ell = s\} \left[(1 - \theta^{\ell-k}) \mathbb{P}\{D^\ell = s\} + \theta^{\ell-k} \mathbb{P}\{Y^\ell = s | Y^k = p\} \right] \\
&\stackrel{(c)}{=} (1 - \theta^{\ell-k}) \mathbb{P}\{D^{\ell+1} = q\} + \theta^{\ell-k} \left[\theta \mathbb{P}\{Y^{\ell+1} = q | Y^k = p\} + (1 - \theta) \mathbb{P}\{D^{\ell+1} = q\} \right] \\
&= (1 - \theta^{\ell+1-k}) \mathbb{P}\{D^{\ell+1} = q\} + \theta^{\ell+1-k} \mathbb{P}\{Y^{\ell+1} = q | Y^k = p\},
\end{aligned}$$

where (c) is by (23). Thus, the claim holds for $\ell + 1 \geq k$ as well. By the claim, $\mathbb{P}\{D^\ell = q | D^k = p\} = (1 - \theta^{\ell-k}) \mathbb{P}\{D^\ell = q\} + \theta^{\ell-k} \mathbb{P}\{Y^\ell = q | Y^k = p\}$, so the lemma follows by arranging the terms. \blacksquare

In the next lemma, we build on Lemma J.1 to bound the gap between the probabilities $\mathbb{P}\{D^k \geq s, D^{k-1} = p | D^\ell = q\}$ and $\mathbb{P}\{D^k \geq s, D^{k-1} = p\}$ for $\ell \geq k + 1$.

Lemma J.2 *For all $p, q, s \in \mathcal{T}$ and $k, \ell \in \mathcal{K}$ with $\ell \geq k + 1$, we have*

$$\sum_{p \in \mathcal{T}} \left| \mathbb{P}\{D^k \geq s, D^{k-1} = p | D^\ell = q\} - \mathbb{P}\{D^k \geq s, D^{k-1} = p\} \right| \leq \frac{1}{\epsilon} (1 - \epsilon)^{\ell-k}.$$

Proof: Using the fact that $\mathbb{P}\{D^{k+1} = q | D^k = p\} \geq \epsilon$ for all $p, q \in \mathcal{T}$ and $k \in \mathcal{K}$, it is simple to check that $\mathbb{P}\{D^k = q\} \geq \epsilon$ for all $q \in \mathcal{T}$ and $k \in \mathcal{K}$. Also, for $\ell \geq k + 1$, by the Bayes rule, we have

$$\begin{aligned}
\frac{\mathbb{P}\{D^k = p, D^{k-1} = r\}}{\mathbb{P}\{D^\ell = q\}} \mathbb{P}\{D^\ell = q | D^k = p\} &= \frac{\mathbb{P}\{D^k = p, D^{k-1} = r\}}{\mathbb{P}\{D^\ell = q\}} \frac{\mathbb{P}\{D^\ell = q, D^k = p\}}{\mathbb{P}\{D^k = p\}} \\
&= \mathbb{P}\{D^{k-1} = r | D^k = p\} \mathbb{P}\{D^k = p | D^\ell = q\} = \mathbb{P}\{D^k = p, D^{k-1} = r | D^\ell = q\}, \tag{24}
\end{aligned}$$

where the last equality uses the fact that given D^k , D^{k-1} is independent of D^ℓ . Noting that we have $\mathbb{P}\{Y^\ell = q | Y^k = p\} - \mathbb{P}\{D^\ell = q\} \in [-1, 1]$, by Lemma J.1, we obtain the chain of inequalities

$-(1-\epsilon)^{\ell-k} \leq \mathbb{P}\{D^\ell = q \mid D^k = p\} - \mathbb{P}\{D^\ell = q\} \leq (1-\epsilon)^{\ell-k}$. In this case, multiplying this chain of inequalities with $\frac{\mathbb{P}\{D^k=p, D^{k-1}=r\}}{\mathbb{P}\{D^\ell=q\}}$, as well as using (24), we have

$$\begin{aligned} & -(1-\epsilon)^{\ell-k} \frac{\mathbb{P}\{D^k=p, D^{k-1}=r\}}{\mathbb{P}\{D^\ell=q\}} \\ & \leq \mathbb{P}\{D^k=p, D^{k-1}=r \mid D^\ell=q\} - \mathbb{P}\{D^k=p, D^{k-1}=r\} \\ & \leq (1-\epsilon)^{\ell-k} \frac{\mathbb{P}\{D^k=p, D^{k-1}=r\}}{\mathbb{P}\{D^\ell=q\}}. \end{aligned}$$

Because $\mathbb{P}\{D^\ell = q\} \geq \epsilon$, the chain of inequalities above yields $-\frac{1}{\epsilon}(1-\epsilon)^{\ell-k} \mathbb{P}\{D^k=p, D^{k-1}=r\} \leq \mathbb{P}\{D^k=p, D^{k-1}=r \mid D^\ell=q\} - \mathbb{P}\{D^k=p, D^{k-1}=r\} \leq \frac{1}{\epsilon}(1-\epsilon)^{\ell-k} \mathbb{P}\{D^k=p, D^{k-1}=r\}$.

To conclude the proof, if we add the last chain of inequalities in the previous paragraph over all $p \geq s$, then we obtain the chain of inequalities

$$\begin{aligned} & -\frac{1}{\epsilon}(1-\epsilon)^{\ell-k} \mathbb{P}\{D^k \geq s, D^{k-1}=r\} \\ & \leq \mathbb{P}\{D^k \geq s, D^{k-1}=r \mid D^\ell=q\} - \mathbb{P}\{D^k \geq s, D^{k-1}=r\} \\ & \leq \frac{1}{\epsilon}(1-\epsilon)^{\ell-k} \mathbb{P}\{D^k \geq s, D^{k-1}=r\}. \end{aligned}$$

Using the fact that $\sum_{r \in \mathcal{T}} \mathbb{P}\{D^k \geq s, D^{k-1}=r\} = \mathbb{P}\{D^k \geq s\} \leq 1$, the chain of inequalities above yields $\sum_{r \in \mathcal{T}} |\mathbb{P}\{D^k \geq s, D^{k-1}=r \mid D^\ell=q\} - \mathbb{P}\{D^k \geq s, D^{k-1}=r\}| \leq \frac{1}{\epsilon}(1-\epsilon)^{\ell-k}$. \blacksquare

Using the lemma above, we can show the first result that we are interested in. In the next lemma, we show that $|M_i^k(\ell) - M_i^k(\ell+1)| \leq \frac{2}{\epsilon^3}$ with probability one.

Lemma J.3 *For all $i \in \mathcal{L}$ and $k, \ell \in \mathcal{K}$ and $\ell+1 \leq k$, with probability one, we have*

$$|M_i^k(\ell) - M_i^k(\ell+1)| \leq \frac{2}{\epsilon^3}.$$

Proof: By the definition of $\Psi_t^k(q)$, we have $\mathbb{E}\{\Psi_t^k(q) \mid \mathbf{D}^{[v,\ell]}\} = \mathbb{P}\{D^k \geq t, D^{k-1}=q \mid \mathbf{D}^{[v,\ell]}\}$ for any $v \in \mathcal{K}$. Noting that $V_i^k = \sum_{v=1}^k \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \Psi_s^v(p) n_{is}^v(p)$, we obtain

$$\begin{aligned} & |M_i^k(\ell) - M_i^k(\ell+1)| = |\mathbb{E}\{V_i^k \mid \mathbf{D}^{[\ell,k]}\} - \mathbb{E}\{V_i^k \mid \mathbf{D}^{[\ell+1,k]}\}| \\ & \leq \sum_{v=1}^k \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} n_{is}^v(p) \left| \mathbb{P}\{D^v \geq s, D^{v-1}=p \mid \mathbf{D}^{[\ell,k]}\} - \mathbb{P}\{D^v \geq s, D^{v-1}=p \mid \mathbf{D}^{[\ell+1,k]}\} \right| \\ & \stackrel{(a)}{=} \sum_{v=1}^{\ell+1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} n_{is}^v(p) \left| \mathbb{P}\{D^v \geq s, D^{v-1}=p \mid \mathbf{D}^{[\ell,k]}\} - \mathbb{P}\{D^v \geq s, D^{v-1}=p \mid \mathbf{D}^{[\ell+1,k]}\} \right|, \quad (25) \end{aligned}$$

where (a) follows from the fact that if $v \geq \ell+2$, then both D^v and D^{v-1} are deterministic functions of $\mathbf{D}^{[\ell+1,k]}$. Therefore, if $v \geq \ell+2$, then the probabilities $\mathbb{P}\{D^v \geq s, D^{v-1}=p \mid \mathbf{D}^{[\ell,k]}\}$ and $\mathbb{P}\{D^v \geq s, D^{v-1}=p \mid \mathbf{D}^{[\ell+1,k]}\}$ take the same value of zero or one. For $v \leq \ell$, given D^ℓ , both

D^v and D^{v-1} are independent of $D^{\ell+1}, \dots, D^k$. Thus, if $v \leq \ell$, then $\mathbb{P}\{D^v \geq s, D^{v-1} = p \mid \mathbf{D}^{[\ell, k]}\} = \mathbb{P}\{D^v \geq s, D^{v-1} = p \mid D^\ell\}$. Similarly, for $v \leq \ell$, given $D^{\ell+1}$, both D^v and D^{v-1} are independent of $D^{\ell+2}, \dots, D^k$. Thus, if $v \leq \ell$, then $\mathbb{P}\{D^v \geq s, D^{v-1} = p \mid \mathbf{D}^{[\ell+1, k]}\} = \mathbb{P}\{D^v \geq s, D^{v-1} = p \mid D^{\ell+1}\}$. Lastly, by the definition of $n_{it}^k(q)$, we have $n_{it}^k(q) = \mathbb{E}\{N_{it}^k(q)\} = \sum_{j \in \mathcal{J}} a_{ij} \gamma \bar{x}_{jt}^k(q) \leq \gamma \sum_{j \in \mathcal{J}} \lambda_{jt}^k \leq 1$, where the first inequality uses the fact that \bar{x} satisfies the second constraint in problem (2). In this case, we can upper bound the expression on the right side of (25) as

$$\begin{aligned}
& \sum_{v=1}^{\ell+1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} n_{is}^v(p) \left| \mathbb{P}\{D^v \geq s, D^{v-1} = p \mid \mathbf{D}^{[\ell, k]}\} - \mathbb{P}\{D^v \geq s, D^{v-1} = p \mid \mathbf{D}^{[\ell+1, k]}\} \right| \\
&= \sum_{v=1}^{\ell} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} n_{is}^v(p) \left| \mathbb{P}\{D^v \geq s, D^{v-1} = p \mid D^\ell\} - \mathbb{P}\{D^v \geq s, D^{v-1} = p \mid D^{\ell+1}\} \right| \\
&\quad + \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} n_{is}^{\ell+1}(p) \left| \mathbb{P}\{D^{\ell+1} \geq s, D^\ell = p \mid \mathbf{D}^{[\ell, k]}\} - \mathbb{P}\{D^{\ell+1} \geq s, D^\ell = p \mid \mathbf{D}^{[\ell+1, k]}\} \right| \\
&\stackrel{(b)}{\leq} \sum_{v=1}^{\ell} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \left| \mathbb{P}\{D^v \geq s, D^{v-1} = p \mid D^\ell\} - \mathbb{P}\{D^v \geq s, D^{v-1} = p\} \right| \\
&\quad + \sum_{v=1}^{\ell} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \left| \mathbb{P}\{D^v \geq s, D^{v-1} = p \mid D^{\ell+1}\} - \mathbb{P}\{D^v \geq s, D^{v-1} = p\} \right| + T^2 \\
&\stackrel{(c)}{\leq} \frac{1}{\epsilon} T \sum_{v=1}^{\ell} (1-\epsilon)^{\ell-v} + \frac{1}{\epsilon} T \sum_{v=1}^{\ell} (1-\epsilon)^{\ell+1-v} + T^2 \\
&\leq \frac{1}{\epsilon} T \left[\frac{1}{\epsilon} + \frac{1-\epsilon}{\epsilon} \right] + T^2, \tag{26}
\end{aligned}$$

where (b) holds by noting that $n_{it}^k(q) \in [0, 1]$ and $|a - b| \leq |a - c| + |c - b|$ in the first sum on the left side of the inequality, as well as bounding the second sum on the left side of the inequality by T^2 , whereas (c) uses Lemma J.2. Because $1 = \sum_{p \in \mathcal{T}} \mathbb{P}\{D^{k+1} = p \mid D^k = q\} \geq T\epsilon$, we get $T \leq 1/\epsilon$, so $\frac{1}{\epsilon} T \left[\frac{1}{\epsilon} + \frac{1-\epsilon}{\epsilon} \right] + T^2 \leq \frac{1}{\epsilon^2} \left[\frac{1}{\epsilon} + \frac{1-\epsilon}{\epsilon} \right] + \frac{1}{\epsilon^2} = \frac{2}{\epsilon^3}$. Thus, the result follows from (25) and (26). \blacksquare

In the next lemma, which is the second result that we are interested in, we give an upper bound on the moment generating function of a Bernoulli random variable.

Lemma J.4 *If Z is Bernoulli with mean μ , then we have $\mathbb{E}\{e^{\lambda(Z-\mu)}\} \leq e^{\mu\lambda^2}$ for all $\lambda \in [0, 1]$.*

Proof: Because e^x is convex in x , over the interval $[0, 1]$, the function e^x lies below the line segment that connects the points $(0, 1)$ and $(1, e)$. Thus, we have $e^x \leq 1 + x(e - 1)$ for all $x \in [0, 1]$. In this case, noting that $e^\lambda = \int_0^\lambda e^x dx + 1$, for all $\lambda \in [0, 1]$, we get $e^\lambda \leq \int_0^\lambda (1 + x(e - 1)) dx + 1 = \lambda + \frac{1}{2}(e - 1)\lambda^2 + 1 \leq 1 + \lambda + \lambda^2$, yielding $e^\lambda - 1 - \lambda \leq \lambda^2$ for all $\lambda \in [0, 1]$. Because Z is Bernoulli with mean μ , we have $\mathbb{P}\{Z = 1\} = \mu$, so $\mathbb{E}\{e^{\lambda Z}\} = (1 - \mu) + \mu e^\lambda = 1 + \mu(e^\lambda - 1) \leq e^{\mu(e^\lambda - 1)}$, where the last inequality holds because $1 + x \leq e^x$ for all $x \in \mathbb{R}$. By the last chain of inequalities, we obtain $\mathbb{E}\{e^{\lambda(Z-\mu)}\} \leq e^{\mu(e^\lambda - 1 - \lambda)} \leq e^{\mu\lambda^2}$, where we use the fact that $e^\lambda - 1 - \lambda \leq \lambda^2$. \blacksquare

Appendix K: Choice of the Tuning Parameters

In the proof of Theorem 4.1, our specific choice of the values for γ and λ is driven by the form of the availability probability bound in Lemma 5.3. To give intuition into these choices, we consider the case with one stage in the selling horizon so that $K = 1$. In this case, the probability in Lemma 5.3 is upper bounded by $e^{c_i \lambda^2 - (1-\gamma)c_i \lambda}$. To make this bound as tight as possible, we find the value of λ that minimizes $c_i \lambda^2 - (1-\gamma)c_i \lambda$, which is given by $\frac{1-\gamma}{2}$. Fixing $\lambda = \frac{1-\gamma}{2}$ and noting that $c_i \geq c_{\min}$, we get $c_i \lambda^2 - (1-\gamma)c_i \lambda = -\frac{1}{4}c_i(1-\gamma)^2 \leq -\frac{1}{4}c_{\min}(1-\gamma)^2$, so the probability in Lemma 5.3 is upper bounded by $e^{-\frac{1}{4}c_{\min}(1-\gamma)^2}$. In this case, by (4), the probability $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\}$ is lower bounded by $1 - L e^{-\frac{1}{4}c_{\min}(1-\gamma)^2}$. By the discussion just after (3), if $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} \geq \alpha$, then we have $\frac{\text{APX}}{\text{ZLP}} \geq \gamma \alpha$. Therefore, we get $\frac{\text{APX}}{\text{ZLP}} \geq \gamma(1 - L e^{-\frac{1}{4}c_{\min}(1-\gamma)^2})$, which implies that we can use $\gamma(1 - L e^{-\frac{1}{4}c_{\min}(1-\gamma)^2})$ as a performance guarantee for our approximate policy. To obtain the best performance guarantee, we can find the value of γ that maximizes $\gamma(1 - L e^{-\frac{1}{4}c_{\min}(1-\gamma)^2})$, but finding the value of γ that maximizes $\gamma(1 - L e^{-\frac{1}{4}c_{\min}(1-\gamma)^2})$ explicitly is difficult.

To get around this difficulty, we make two observations. First, we parameterize the value of γ as $\gamma = 1 - 2c_{\min}^\alpha$ for some choice of $\alpha \in \mathbb{R}$. In this case, the performance guarantee for our approximate policy takes the form $\gamma(1 - L e^{-\frac{1}{4}c_{\min}(1-\gamma)^2}) = (1 - 2c_{\min}^\alpha)(1 - L e^{-c_{\min}^{2\alpha+1}})$. As c_{\min} becomes arbitrarily large, we would like the performance guarantee to be arbitrarily close to one. If $\alpha \in (-\frac{1}{2}, 0)$, then we have $\lim_{c_{\min} \rightarrow \infty} (1 - 2c_{\min}^\alpha)(1 - L e^{-c_{\min}^{2\alpha+1}}) = 1$. Thus, we choose the value of α in the interval $(-\frac{1}{2}, 0)$. Second, because $c_{\min} \geq 1$ and $\alpha > -\frac{1}{2}$, for large enough values of c_{\min} , we have $(1 - 2c_{\min}^\alpha)(1 - L e^{-c_{\min}^{2\alpha+1}}) \leq 1 - 2c_{\min}^\alpha \leq 1 - 2c_{\min}^{-\frac{1}{2}}$, so the best possible performance guarantee that we can hope to obtain is $1 - 2\frac{1}{\sqrt{c_{\min}}}$. If we fix $\alpha = \frac{\log \log c_{\min}}{2 \log c_{\min}} - \frac{1}{2}$, then $2\alpha + 1 = \frac{\log \log c_{\min}}{\log c_{\min}}$, so using the identity $x^{(\log \log x)/(2 \log x)} = \sqrt{\log x}$, the performance guarantee for the approximate policy becomes $(1 - 2c_{\min}^\alpha)(1 - L e^{-c_{\min}^{2\alpha+1}}) = \left(1 - 2\sqrt{\frac{\log c_{\min}}{c_{\min}}}\right) \left(1 - \frac{L}{c_{\min}}\right)$. Therefore, our choice of γ ensures that our approximate policy gets to have a performance guarantee of $\left(1 - 2\sqrt{\frac{\log c_{\min}}{c_{\min}}}\right) \left(1 - \frac{L}{c_{\min}}\right)$, whereas the best possible performance guarantee that we can hope for is $1 - O\left(\frac{1}{\sqrt{c_{\min}}}\right)$.

Note that $\left(1 - 2\sqrt{\frac{\log c_{\min}}{c_{\min}}}\right) \left(1 - \frac{L}{c_{\min}}\right) \geq 1 - 2\sqrt{\frac{\log c_{\min}}{c_{\min}}} - \frac{L}{c_{\min}} = 1 - O\left(\sqrt{\frac{\log c_{\min}}{c_{\min}}}\right)$, so our choice of γ ensures that our approximate policy has a performance guarantee of $1 - O\left(\sqrt{\frac{\log c_{\min}}{c_{\min}}}\right)$ and the best possible performance guarantee that we can hope for is $1 - O\left(\frac{1}{\sqrt{c_{\min}}}\right)$. Therefore, our choice of γ ensures that we match the best possible performance guarantee up to a logarithmic factor. Using the identity $x^{(\log \log x)/(2 \log x)} = \sqrt{\log x}$, if we fix $\alpha = \frac{\log \log c_{\min}}{2 \log c_{\min}} - \frac{1}{2}$, then our choice of γ is given by $\gamma = 1 - 2c_{\min}^\alpha = 1 - 2\sqrt{\frac{\log c_{\min}}{c_{\min}}} = 1 - \frac{\sqrt{4c_{\min} \log c_{\min}}}{c_{\min}}$, which matches the choice of γ in the proof of Theorem 4.1 when there is one stage in the selling horizon so that $K = 1$. The preceding discussion sheds light into our choice of the tuning parameter γ for the case with one stage in the selling

horizon. We can use precisely the same reasoning for the case with multiple stages in the selling horizon, but the calculations with multiple stages get slightly more tedious.

Appendix L: Independent Demands in Different Stages

We study performance guarantees for an approximate policy from our fluid approximation when the demand random variables D^1, \dots, D^K in different stages are independent of each other.

L.1 Approximate Policy and Asymptotic Optimality

In this section, we elaborate on the fact that having K stages with the demand random variables D^1, \dots, D^K being independent of each other is not equivalent to having one stage with the number of customer arrivals being given by the random variable $D^1 + \dots + D^K$. Next, when working with independent demand random variables, we assume that the demands in different stages are sub-Gaussian. We explain that this assumption is milder than assuming that $\mathbb{P}\{D^{k+1} = p\} \geq \epsilon$ for all $p \in \mathcal{T}$ and $k \in \mathcal{K}$ for some $\epsilon > 0$, which would have been the analogue of the assumption in Section 2 when the demand random variables in different stages are independent of each other. We carefully contrast these two assumptions. Following this discussion, we provide an approximate policy that uses our fluid approximation under independent demands. We give a performance guarantee for this policy under sub-Gaussian demand random variables that are independent of each other. Lastly, we demonstrate that our fluid approximation has a simplified form under independent demand random variables and other intuitive fluid approximations that make seemingly harmless modifications in our fluid approximation do not yield asymptotically tight upper bounds.

Having K stages with the demand random variables D^1, \dots, D^K being independent of each other is not equivalent to having one stage with the number of customer arrivals being given by the random variable $D^1 + \dots + D^K$. In particular, the resolution of the demand uncertainty in the two cases are different. Consider having K stages with the demand random variables D^1, \dots, D^K being independent of each other. If we are in stage k and there have been t customer arrivals in the current stage, then the probability of having one more customer arrival is $\mathbb{P}\{D^k \geq t + 1 \mid D^k \geq t\} + \mathbb{P}\{D^k = t \mid D^k \geq t\} \mathbb{P}\{D^{k+1} + \dots + D^K \geq 1\}$, which corresponds to the probability of having one more customer arrival in the current stage plus the probability of having no more customer arrivals in the current stage and one more customer arrival in the remaining stages. If the stages correspond to, for example, weeks, then this first demand model is aware of the calendar in the sense that it distinguishes between having one more customer arrival in the current week and finishing the current week with no more customer arrivals. On the other hand, consider having one stage with the number of customer arrivals being given by the random variable $D^1 + \dots + D^K$. If there

have been v customer arrivals so far, then the probability of having one more customer arrival is $\mathbb{P}\{D^1 + \dots + D^K \geq v + 1 \mid D^1 + \dots + D^K \geq v\}$, which does not explicitly consider the distribution of the random demand over different stages. If the stages correspond to weeks, then this second demand model is not aware of the calendar in the sense that it does not pay attention to the beginning and end of each week. To our knowledge, demand models with uncertainty in the demand resolving sequentially over multiple stages have not been considered. Thus, even when the demands in different stages are independent, our fluid approximation yields new results.

When the demand random variables in different stages are independent, we drop the assumption that $\mathbb{P}\{D^{k+1} = p \mid D^k = q\} \geq \epsilon$ for all $p, q \in \mathcal{T}$ and $k \in \mathcal{K}$ for some $\epsilon > 0$. Instead, we assume that the demand random variable in each stage is sub-Gaussian, so we have $\mathbb{E}\{e^{\lambda(D^k - \mathbb{E}\{D^k\})}\} \leq e^{\frac{\sigma^2}{200}\lambda^2}$ for all $k \in \mathcal{K}$ and $\lambda \geq 0$ for some $\sigma^2 > 0$. The assumption of sub-Gaussian demand is mild, as the class of sub-Gaussian random variables is large; see, for example, Section 2.1.2 in Wainwright (2019). The parameter $\frac{\sigma^2}{200}$ is known as the variance proxy because if D^k satisfies the sub-Gaussian assumption, then the variance of D^k is at most $\frac{\sigma^2}{200}$. We scale the variance proxy by $1/200$ for notational uniformity in our proofs. Under independent demands, we continue using the fluid approximation in (2), but this linear program admits obvious simplifications under independent demands. For example, we can write the probability $\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\}$ in the first constraint as $\mathbb{P}\{D^\ell \geq s\} \mathbb{P}\{D^{\ell-1} = p\}$ for $\ell < k - 1$. To construct an approximate policy from the fluid approximation, we solve problem (2) once at the beginning of the selling horizon. Letting $\bar{x} = (\bar{x}_{jt}^k(q) : j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$ be an optimal solution, we make the decisions as follows.

Approximate Policy under Independent Demands:

Using $\gamma \in [0, 1]$ to denote a tuning parameter, if we have a request for product j at time period t in stage k , then we are willing to accept the request with probability $\gamma \sum_{q \in \mathcal{T}} \mathbb{P}\{D^{k-1} = q\} \frac{\bar{x}_{jt}^k(q)}{\lambda_{jt}^k}$. If we are willing to accept the request and there are enough resource capacities to accept the request, then we accept the request. Otherwise, we reject.

In the next theorem, we give a performance guarantee for the approximate policy. We continue using APX to denote the total expected revenue of the approximate policy.

Theorem L.1 (Independent Demands) *Under independent demands, there exists a choice of the tuning parameter γ such that we have*

$$\frac{\text{APX}}{\text{OPT}} \geq \frac{\text{APX}}{\overline{Z}_{\text{LP}}} \geq \max \left\{ \frac{1}{4L}, \left(1 - 4 \frac{\sqrt{(c_{\min} + \sigma^2(K-1)) \log c_{\min}}}{c_{\min}} - \frac{L}{c_{\min}} \right) \right\}.$$

We give the proof of the theorem shortly in Appendix L.2. Although the demand random variables in the theorem above are independent, we are unable to use concentration inequalities

for sums of independent random variables in the proof. In particular, using N_{it}^k to capture the Bernoulli random variable that takes value one if there is a request for a product at time period t in stage k that uses resource i and the approximate policy is willing to accept the request, we can use $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \mathbf{1}(D^\ell \geq s) N_{is}^\ell + \sum_{s=1}^t \mathbf{1}(D^k \geq s) N_{is}^k$ to upper bound the capacity consumption of resource i up to and including time period t in stage k . This random variable is the analogue of $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{q \in \mathcal{T}} \Psi_s^\ell(q) N_{is}^\ell(q) + \sum_{s=1}^t \sum_{q \in \mathcal{T}} \Psi_s^k(q) N_{is}^k(q)$ used right after (3) in Section 5. Considering the sum $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \mathbf{1}(D^\ell \geq s) N_{is}^\ell + \sum_{s=1}^t \mathbf{1}(D^k \geq s) N_{is}^k$, even if the demands are independent, for fixed stage ℓ , the terms $\{\mathbf{1}(D^\ell \geq s) N_{is}^\ell : s \in \mathcal{T}\}$ in this sum are not independent due to the random variable $\mathbf{1}(D^\ell \geq s)$. Thus, we cannot bound the tail probabilities of the capacity consumptions by using concentration inequalities for sums of independent random variables. We derive our concentration inequalities through moment generating function bounds.

When the demands in different stages are dependent, to establish our performance guarantee, we assume that $\mathbb{P}\{D^k = t\} \geq \epsilon$ for all $t \in \mathcal{T}$ and $k \in \mathcal{K}$ for some $\epsilon > 0$. If $\mathbb{P}\{D^k = t\} \geq \epsilon$ for all $t \in \mathcal{T}$, then $D^k \leq 1/\epsilon$ with probability one. When the demands in different stages are independent, to establish our performance guarantee, we assume that the demands are sub-Gaussian. By Lemma 5.1 in Dubhashi and Panconesi (2009), bounded random variables are sub-Gaussian, so the latter assumption is milder, but the milder assumption comes at the expense of having to assume independence between the demands. The interpretation Theorem L.1 is similar to that of Theorem 4.1. The approximate policy has a constant-factor performance guarantee, as long as L is fixed. Furthermore, if we scale both the number of stages and resource capacities with the same rate θ , then the relative gap between the total expected revenue of the approximate policy and the optimal total expected revenue converges to one with rate $1 - \frac{1}{\sqrt{\theta}}$.

There are important differences in the analyses of our approximate policies under dependent and independent demands. Under independent demands, we construct concentration inequalities for a random variable of the form $\sum_{k=1}^K \sum_{t=1}^{D^k} B_t^k$, where $\{B_t^k : t \in \mathcal{T}, k \in \mathcal{K}\}$ are independent Bernoulli random variables. The main difficulty in constructing these concentration inequalities is that the quantity D^k is random. Thus, we deal with sums of random numbers of Bernoulli random variables, while the Bernoulli random variables in the sum being independent of each other. On the other hand, under dependent demands, we construct concentration inequalities for a random variable of the form $\sum_{k=1}^K \sum_{t=1}^{D^k} B_t^k(D^{k-1})$, where $\{B_t^k(D^{k-1}) : t \in \mathcal{T}, k \in \mathcal{K}\}$ are dependent Bernoulli random variables, because the distribution of $B_t^k(D^{k-1})$ depends on D^{k-1} and $\{D^1, \dots, D^K\}$ are dependent on each other. Thus, we deal with sums of random numbers of Bernoulli random variables, while the Bernoulli random variables in the sum being dependent on each other.

L.2 Performance Guarantee

We give a proof for Theorem L.1. Our starting point is similar to the outline of Section 5, but we will construct different concentration inequalities to exploit the sub-Gaussian property.

Preliminary Random Variables and Availability Probabilities:

We define four classes of Bernoulli random variables for each $k \in \mathcal{K}$ and $t \in \mathcal{T}$. We worked with the analogues of these random variables in Section 5.

- *Demand in Each Stage.* The random variable Ψ_t^k takes value one if we reach time period t in stage k before this stage is over. In other words, recalling that we use $\mathbf{1}(\cdot)$ to denote the indicator function, we simply have $\Psi_t^k = \mathbf{1}(D^k \geq t)$.

- *Product Request.* For each $j \in \mathcal{J}$, the random variable A_{jt}^k takes value one if the customer arriving at time period t in stage k requests product j . We have $\mathbb{P}\{A_{jt}^k = 1\} = \lambda_{jt}^k$. The random variables $\{A_{jt}^k : t \in \mathcal{T}, k \in \mathcal{K}\}$ are independent of each other.

- *Policy Decision.* For each $j \in \mathcal{J}$, the random variable X_{jt}^k takes value one if the approximate policy is willing to accept a request for product j at time period in stage k . By the definition of the approximate policy under independent demands, $\mathbb{P}\{X_{jt}^k = 1\} = \gamma \sum_{q \in \mathcal{T}} \mathbb{P}\{D^{k-1} = q\} \frac{\bar{x}_{jt}^k(q)}{\lambda_{jt}^k}$.

- *Availability.* For each $j \in \mathcal{J}$, the random variable G_{jt}^k takes value one if we have enough capacity to accept a request for product j at time period t in stage k under the approximate policy. Instead of calculating the probability $\mathbb{P}\{G_{jt}^k = 1\}$, we will lower bound $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k = 1\}$.

Under independent demands, we drop the assumption that $\mathbb{P}\{D^{k+1} = p \mid D^k = q\} \geq \epsilon$ for all $p, q \in \mathcal{T}$ and $k \in \mathcal{K}$ for some $\epsilon > 0$. Instead, we assume that $\mathbb{E}\{e^{\lambda(D^k - \mathbb{E}\{D^k\})}\} \leq e^{\frac{\sigma^2}{200}\lambda^2}$ for all $k \in \mathcal{K}$ and $\lambda \geq 0$ for some $\sigma^2 > 0$. Throughout this section, when we refer to the approximate policy, we mean the approximate policy under independent demands as described in Section L. Under the approximate policy, the sales for product j at time period t in stage k is given by the random variable $\Psi_t^k G_{jt}^k A_{jt}^k X_{jt}^k$. Taking expectations, the expected sales for product j at time period t in stage k is $\mathbb{P}\{\Psi_t^k = 1\} \mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k = 1\} \mathbb{P}\{A_{jt}^k = 1\} \mathbb{P}\{X_{jt}^k = 1\}$, where we use the fact that the random variables A_{jt}^k and X_{jt}^k are independent of the demands and remaining capacities of the resources. Noting that $\mathbb{P}\{A_{jt}^k = 1\} = \lambda_{jt}^k$ and $\mathbb{P}\{X_{jt}^k = 1\} = \gamma \sum_{q \in \mathcal{T}} \mathbb{P}\{D^{k-1} = q\} \frac{\bar{x}_{jt}^k(q)}{\lambda_{jt}^k}$, the last expectation is given by $\mathbb{P}\{D^k \geq t\} \mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k = 1\} \sum_{q \in \mathcal{T}} \gamma \mathbb{P}\{D^{k-1} = q\} \bar{x}_{jt}^k(q)$. Therefore, the total expected revenue of the approximate policy is

$$\text{APX} = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \sum_{q \in \mathcal{T}} f_j \mathbb{P}\{D^k \geq t\} \mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k = 1\} \gamma \mathbb{P}\{D^{k-1} = q\} \bar{x}_{jt}^k(q).$$

Because \bar{x} is an optimal solution to problem (2), the optimal objective value of problem (2) is $\bar{Z}_{\text{LP}} = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \sum_{q \in \mathcal{T}} f_j \mathbb{P}\{D^k \geq t\} \mathbb{P}\{D^{k-1} = q\} \bar{x}_{jt}^k(q)$, where we use the fact that the

demands in different stages are independent, so $\mathbb{P}\{D^k \geq t, D^{k-1} = q\} = \mathbb{P}\{D^k \geq t\} \mathbb{P}\{D^{k-1} = q\}$. Thus, comparing the expressions for the total expected revenue APX of the approximate policy and the optimal objective value \bar{Z}_{LP} of the linear program in (2), it follows that if we can show that $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k = 1\} \geq \alpha$, then we obtain $\text{APX} \geq \gamma \alpha \bar{Z}_{\text{LP}}$. We focus on lower bounding the availability probability $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k = 1\}$. Under the approximate policy, the capacity consumption of resource i at time period t in stage k is given by the random variable $\sum_{j \in \mathcal{J}} a_{ij} \Psi_t^k G_{jt}^k A_{jt}^k X_{jt}^k$, which implies that the capacity consumption of resource i at time period t in stage k is upper bounded by $\sum_{j \in \mathcal{J}} a_{ij} \Psi_t^k A_{jt}^k X_{jt}^k$. Letting $N_{it}^k = \sum_{j \in \mathcal{J}} a_{ij} A_{jt}^k X_{jt}^k$, we write the upper bound on the capacity consumption of resource i at time period t in stage k as $\Psi_t^k N_{it}^k$. In this case, if we have $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \Psi_s^\ell N_{is}^\ell + \sum_{s=1}^t \Psi_s^k N_{is}^k < c_i$, then the total capacity consumption of resource i up to and including time period t in stage k does not exceed the capacity of the resource, in which case, we must have capacity available for resource i at time period t in stage k . Therefore, using the same line of reasoning that we followed to obtain the chain of inequalities in (4), we can lower bound the availability probability as

$$\mathbb{P}\{G_{jt} = 1 \mid \Psi_t^k = 1\} \geq 1 - \sum_{i \in \mathcal{L}_j} \mathbb{P}\left\{ \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \Psi_s^\ell N_{is}^\ell + \sum_{s=1}^t \Psi_s^k N_{is}^k \geq c_i \mid D^k \geq t \right\}. \quad (27)$$

The discussion so far closely followed the one at the beginning of Section 5, but we need to deviate from that discussion to exploit the sub-Gaussian assumption.

Moment Generating Function Bounds:

Note that N_{it}^k is a Bernoulli random variable. We define $n_{it}^k = \mathbb{E}\{N_{it}^k\}$. In the next lemma, we bound the moment generating function of the squared deviation of the demand around its mean.

Lemma L.2 *Letting $\mu^k = \mathbb{E}\{D^k\}$, for all $k \in \mathcal{K}$, if $|\lambda| \leq \frac{2}{\sigma}$, then we have*

$$\mathbb{E}\{e^{\lambda^2 (D^k - \mu^k)^2}\} \leq e^{\frac{1}{4} \sigma^2 \lambda^2}.$$

Proof: In an auxiliary lemma, labeled as Lemma L.5, given at the end of this section, we show that if Z is a mean-zero sub-Gaussian random variable with variance proxy M^2 so that $\mathbb{E}\{e^{\lambda Z}\} \leq e^{M^2 \lambda^2}$ for all $\lambda \in \mathbb{R}$, then we have $\mathbb{E}\{e^{\theta^2 Z^2}\} \leq e^{(7M\theta)^2}$ for all $|\theta| \leq \frac{1}{7M}$. Because $D^k - \mu^k$ is a mean-zero sub-Gaussian random variable with variance proxy $\frac{\sigma^2}{200}$, using Lemma L.5 with $M = \frac{\sigma}{\sqrt{200}}$, we have $\mathbb{E}\{e^{\theta^2 (D^k - \mu^k)^2}\} \leq e^{(7\frac{\sigma}{\sqrt{200}}\theta)^2}$ for all $|\theta| \leq \frac{1}{7\sigma/\sqrt{200}}$. We have $\frac{1}{7\sigma/\sqrt{200}} \geq \frac{2}{\sigma}$. Thus, if we have $|\theta| \leq \frac{2}{\sigma}$, then we also have $|\theta| \leq \frac{1}{7\sigma/\sqrt{200}}$, in which case, $\mathbb{E}\{e^{\theta^2 (D^k - \mu^k)^2}\} \leq e^{(7\frac{\sigma}{\sqrt{200}}\theta)^2} = e^{\frac{49}{200} \sigma^2 \theta^2} \leq e^{\frac{1}{4} \sigma^2 \theta^2}$. ■

In the next lemma, we use the lemma above to bound the moment generating function of $\sum_{s \in \mathcal{T}} \mathbf{1}(D^k \geq s) n_{is}^k$, which will be key to bounding the availability probabilities.

Lemma L.3 For all $k \in \mathcal{K}$, $i \in \mathcal{L}$ and $\lambda \in \mathbb{R}$, we have

$$\mathbb{E}\{e^{\lambda \sum_{s \in \mathcal{T}} n_{is}^k \mathbf{1}(D^k \geq s)}\} \leq e^{\sigma^2 \lambda^2 + \lambda \sum_{s \in \mathcal{T}} n_{is}^k \mathbb{P}\{D^k \geq s\}}.$$

Proof: For the moment, we assume that $|\lambda| \leq \frac{1}{\sigma}$. We will need two inequalities. First, we use the random variable R^k to denote an independent and identically distributed copy of D^k . Noting that $\sum_{s \in \mathcal{T}} n_{is}^k \mathbf{1}(D^k \geq s) = \sum_{s=1}^{D^k} n_{is}^k$, using the fact that $n_{is}^k \in [0, 1]$, we get the chain of inequalities $|\sum_{s \in \mathcal{T}} n_{is}^k (\mathbf{1}(D^k \geq s) - \mathbf{1}(R^k \geq s))| = |\sum_{s=1}^{D^k} n_{is}^k - \sum_{s=1}^{R^k} n_{is}^k| \leq |D^k - R^k|$. It is simple to verify the inequality $(a - b)^2 \leq 2(a - \delta)^2 + 2(b - \delta)^2$. Thus, letting $\mu^k = \mathbb{E}\{D^k\} = \mathbb{E}\{R^k\}$, we obtain $[\sum_{s \in \mathcal{T}} n_{is}^k (\mathbf{1}(D^k \geq s) - \mathbf{1}(R^k \geq s))]^2 \leq (D^k - R^k)^2 \leq 2(D^k - \mu^k)^2 + 2(R^k - \mu^k)^2$. Second, because $|\lambda| \leq \frac{1}{\sigma}$, we have $|\sqrt{2}\lambda| \leq \frac{\sqrt{2}}{\sigma} \leq \frac{2}{\sigma}$, so by Lemma L.2, $\mathbb{E}\{e^{2\lambda^2(D^k - \mu^k)^2}\} \leq e^{\frac{1}{2}\sigma^2\lambda^2}$. Thus, we have

$$\begin{aligned} \mathbb{E}\{e^{\lambda^2[\sum_{s \in \mathcal{T}} n_{is}^k (\mathbf{1}(D^k \geq s) - \mathbb{P}\{D^k \geq s\})]^2}\} &= \mathbb{E}\{e^{\lambda^2[\sum_{s \in \mathcal{T}} n_{is}^k (\mathbf{1}(D^k \geq s) - \mathbb{E}\{\mathbf{1}(R^k \geq s)\})]^2}\} \\ &\stackrel{(a)}{\leq} \mathbb{E}\{e^{\lambda^2[\sum_{s \in \mathcal{T}} n_{is}^k (\mathbf{1}(D^k \geq s) - \mathbf{1}(R^k \geq s))]^2}\} \stackrel{(b)}{\leq} \mathbb{E}\{e^{2\lambda^2(D^k - \mu^k)^2}\} \mathbb{E}\{e^{2\lambda^2(R^k - \mu^k)^2}\} \stackrel{(c)}{\leq} e^{\sigma^2\lambda^2}, \end{aligned} \quad (28)$$

where (a) follows by the Jensen inequality and the fact that $e^{\lambda^2 x^2}$ is convex in x , (b) uses the first inequality in this paragraph, as well as the fact that the random variables D^k and R^k are independent and (c) uses the second inequality in this paragraph. In an auxiliary lemma, labeled as Lemma L.6, given at the end of this section, we show that if Z is a mean-zero random variable that satisfies $\mathbb{E}\{e^{\lambda^2 Z^2}\} \leq e^{M^2 \lambda^2}$ for all $|\lambda| \leq \frac{1}{M}$, then we have $\mathbb{E}\{e^{\theta Z}\} \leq e^{M^2 \theta^2}$ for all $\theta \in \mathbb{R}$. By (28), we have $\mathbb{E}\{e^{\lambda^2[\sum_{s \in \mathcal{T}} n_{is}^k (\mathbf{1}(D^k \geq s) - \mathbb{P}\{D^k \geq s\})]^2}\} \leq e^{\sigma^2 \lambda^2}$ for all $|\lambda| \leq \frac{1}{\sigma}$. Thus, using Lemma L.6 with $M = \sigma$, $\mathbb{E}\{e^{\theta \sum_{s \in \mathcal{T}} n_{is}^k (\mathbf{1}(D^k \geq s) - \mathbb{P}\{D^k \geq s\})}\} \leq e^{\sigma^2 \theta^2}$ for all $\theta \in \mathbb{R}$, which is the desired result. \blacksquare

Performance Guarantee for the Approximate Policy:

In the next lemma, we use the moment generating function bound in Lemma L.3 to give a lower bound on the availability probabilities on the right side of (27).

Lemma L.4 For all $k \in \mathcal{K}$, $i \in \mathcal{L}$, $t \in \mathcal{T}$ and $\lambda \in [0, 1]$, we have

$$\mathbb{P}\left\{\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \mathbf{1}(D^\ell \geq s) N_{is}^\ell + \sum_{s=1}^t N_{is}^k \geq c_i \mid D^k \geq t\right\} \leq e^{(c_i + 4\sigma^2(k-1))\lambda^2 - (1-\gamma)c_i\lambda}.$$

Proof: Given D^k , $\mathbf{1}(D^k \geq s) N_{is}^k$ is a Bernoulli random variable with mean $\mathbf{1}(D^k \geq s) n_{is}^k$, so by Lemma J.4, we have $\mathbb{E}\{e^{\lambda \mathbf{1}(D^k \geq s) (N_{is}^k - n_{is}^k)} \mid D^k\} \leq e^{\lambda^2 \mathbf{1}(D^k \geq s) n_{is}^k}$. In this case, we obtain

$$\mathbb{E}\{e^{\lambda \sum_{s \in \mathcal{T}} \mathbf{1}(D^k \geq s) N_{is}^k} \mid D^k\} = e^{\lambda \sum_{s \in \mathcal{T}} \mathbf{1}(D^k \geq s) n_{is}^k} \mathbb{E}\{e^{\lambda \sum_{s \in \mathcal{T}} \mathbf{1}(D^k \geq s) (N_{is}^k - n_{is}^k)} \mid D^k\} \leq e^{(\lambda + \lambda^2) \sum_{s \in \mathcal{T}} \mathbf{1}(D^k \geq s) n_{is}^k},$$

where we use the fact that $\{N_{is}^k : s \in \mathcal{T}\}$ are independent of each other. Taking expectations in the chain of inequalities above, we obtain $\mathbb{E}\{e^{\lambda \sum_{s \in \mathcal{T}} \mathbf{1}(D^k \geq s) N_{is}^k}\} \leq \mathbb{E}\{e^{(\lambda + \lambda^2) \sum_{s \in \mathcal{T}} \mathbf{1}(D^k \geq s) n_{is}^k}\}$. Using

the same argument, we have $\mathbb{E}\{e^{\lambda \sum_{s=1}^t N_{is}^k}\} \leq e^{(\lambda^2 + \lambda) \sum_{s=1}^t n_{is}^k}$. On the other hand, by the discussion in Section 3, we can replace the probability $\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\}$ in the first constraint in problem (2) with $\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p \mid D^{k-1} = q\}$. Furthermore, \bar{x} is a feasible solution to problem (2), so it satisfies the first constraint. Multiplying this constraint with $\mathbb{P}\{D^{k-1} = q\}$ and adding over all $q \in \mathcal{T}$, we get $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p\} \bar{x}_{js}^\ell(p) + \sum_{s=1}^t \sum_{j \in \mathcal{J}} \sum_{q \in \mathcal{T}} a_{ij} \mathbb{P}\{D^{k-1} = q\} \bar{x}_{js}^k(q) \leq c_i$. By the definition of N_{it}^k , we also have the identity $n_{it}^k = \mathbb{E}\{N_{it}^k\} = \gamma \sum_{j \in \mathcal{J}} a_{ij} \sum_{q \in \mathcal{T}} \mathbb{P}\{D^{k-1} = q\} \bar{x}_{jt}^k(q)$. Thus, noting that the demands in different stages are independent of each other, so that $\mathbb{P}\{D^\ell \geq s, D^{\ell-1} = p\} = \mathbb{P}\{D^\ell \geq s\} \mathbb{P}\{D^{\ell-1} = p\}$, the last inequality yields $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \mathbb{P}\{D^\ell \geq s\} n_{is}^\ell + \sum_{s=1}^t n_{is}^k \leq \gamma c_i$. Noting that the random variables $\{N_{it}^k : t \in \mathcal{T}, k \in \mathcal{K}\}$ and $\{D^k : k \in \mathcal{K}\}$ are all independent of each other, we get

$$\begin{aligned}
\mathbb{P}\left\{\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \mathbf{1}(D^\ell \geq s) N_{is}^\ell + \sum_{s=1}^t N_{is}^k \geq c_i \mid D^k \geq t\right\} &= \mathbb{P}\{e^{\lambda(\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \mathbf{1}(D^\ell \geq s) N_{is}^\ell + \sum_{s=1}^t N_{is}^k)} \geq e^{\lambda c_i}\} \\
&\stackrel{(a)}{\leq} \frac{1}{e^{\lambda c_i}} \mathbb{E}\{e^{\lambda(\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \mathbf{1}(D^\ell \geq s) N_{is}^\ell + \sum_{s=1}^t N_{is}^k)}\} \\
&= \frac{1}{e^{\lambda c_i}} \prod_{\ell=1}^{k-1} \mathbb{E}\{e^{\lambda \sum_{s \in \mathcal{T}} \mathbf{1}(D^\ell \geq s) N_{is}^\ell}\} \mathbb{E}\{e^{\lambda \sum_{s=1}^t N_{is}^k}\} \\
&\stackrel{(b)}{\leq} \frac{1}{e^{\lambda c_i}} \prod_{\ell=1}^{k-1} \mathbb{E}\{e^{(\lambda^2 + \lambda) \sum_{s \in \mathcal{T}} \mathbf{1}(D^\ell \geq s) n_{is}^\ell}\} e^{(\lambda^2 + \lambda) \sum_{s=1}^t n_{is}^k} \\
&\stackrel{(c)}{\leq} \frac{1}{e^{\lambda c_i}} \prod_{\ell=1}^{k-1} e^{4\sigma^2 \lambda^2 + (\lambda^2 + \lambda) \sum_{s \in \mathcal{T}} \mathbb{P}\{D^\ell \geq s\} n_{is}^\ell} e^{(\lambda^2 + \lambda) \sum_{s=1}^t n_{is}^k} \\
&= \frac{1}{e^{\lambda c_i}} e^{4\sigma^2(k-1)\lambda^2} e^{(\lambda^2 + \lambda)(\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \mathbb{P}\{D^\ell \geq s\} n_{is}^\ell + \sum_{s=1}^t n_{is}^k)} \\
&\stackrel{(d)}{\leq} e^{(c_i + 4\sigma^2(k-1))\lambda^2 - (1-\gamma)c_i\lambda},
\end{aligned}$$

where (a) is the Markov inequality, (b) is by the discussion earlier in the proof, (c) uses Lemma L.3 and $4\lambda^2 \geq (\lambda^2 + \lambda)^2$ for $\lambda \in [0, 1]$ and (d) holds by $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \mathbb{P}\{D^\ell \geq s\} n_{is}^\ell + \sum_{s=1}^t n_{is}^k \leq \gamma c_i$. ■

Using specific values for γ and λ in Lemma L.4 will provide a lower bound for the availability probabilities, which, in turn, will yield a performance guarantee for the approximate policy.

Proof of Theorem L.1:

Identifying $4\sigma^2$ with $\frac{3}{e^6}$, the bounds in Lemmas 5.3 and L.4 have the same form. Thus, choosing $\gamma = 1 - \frac{\sqrt{4(c_{\min} + 4\sigma^2(K-1)) \log c_{\min}}}{c_{\min}}$ and $\lambda = \frac{(1-\gamma)c_i}{2(c_i + 4\sigma^2(K-1))}$, following the proof of Theorem 4.1, we get

$$\frac{\text{APX}}{\bar{Z}_{\text{LP}}} \geq \left(1 - 4 \frac{\sqrt{(c_{\min} + \sigma^2(K-1)) \log c_{\min}}}{c_{\min}} - \frac{L}{c_{\min}}\right).$$

The proof of the inequality above precisely follows the argument at the end of Section 5 line by line. On the other hand, we can precisely follow the argument in Appendix I line by line to show

that $\frac{\text{APX}}{\text{ZLP}} \geq \frac{1}{4L}$. In this case, the desired result follows by collecting these two inequalities together, as well as noting that we have $\frac{\text{APX}}{\text{OPT}} \geq \frac{\text{APX}}{\text{ZLP}}$ by Theorem 3.1. \blacksquare

In the proofs of Lemmas L.2 and L.3, we used two results on sub-Gaussian random variables. We used Lemma L.5 in the proof of Lemma L.2 and Lemma L.6 in the proof of Lemma L.3.

Lemma L.5 *If Z is a mean-zero sub-Gaussian random variable such that $\mathbb{E}\{e^{\lambda Z}\} \leq e^{M^2\lambda^2}$ for all $\lambda \in \mathbb{R}$, then we have $\mathbb{E}\{e^{\theta^2 Z^2}\} \leq e^{(7M\theta)^2}$ for all $|\theta| \leq \frac{1}{7M}$.*

Proof: By the discussion that follows Definition 2.2 in Wainwright (2019), if $\mathbb{E}\{e^{\lambda Z}\} \leq e^{M^2\lambda^2}$ for all $\lambda \in \mathbb{R}$, then $\mathbb{P}\{|Z| \geq t\} \leq 2e^{-\frac{t^2}{4M^2}}$ for all $t \geq 0$. Letting $\|Z\|_{L^p} = (\mathbb{E}\{|Z|^p\})^{1/p}$ for $p \in \mathbb{Z}_{++}$, we get

$$\begin{aligned} \mathbb{E}\{|Z|^p\} &= \int_0^\infty \mathbb{P}\{|Z|^p \geq u\} du \stackrel{(a)}{=} \int_0^\infty \mathbb{P}\{|Z| \geq t\} p t^{p-1} dt \leq 2 \int_0^\infty e^{-\frac{t^2}{4M^2}} p t^{p-1} dt \\ &\stackrel{(b)}{=} 2p \int_0^\infty e^{-z} (2M\sqrt{z})^{p-1} \frac{M}{\sqrt{z}} dz = p(2M)^p \int_0^\infty e^{-z} z^{\frac{p}{2}-1} dz \stackrel{(c)}{=} p(2M)^p \Gamma(p/2) \stackrel{(d)}{\leq} p(2M)^p (p/2)^{p/2}, \end{aligned}$$

where (a) uses the change of variables $u = t^p$, (b) uses the change of variables $z = t^2/(4M^2)$, (c) is the definition of the gamma function and (d) holds because the gamma function satisfies $\Gamma(x) \leq x^x$. By the chain of inequalities above, we obtain $\|Z\|_{L^p} \leq p^{1/p} M\sqrt{2p}$. It is simple to verify that $\max_{p \in \mathbb{Z}_{++}} p^{1/p} = 3^{1/3}$, so the last inequality yields $\|Z\|_{L^p} \leq 3^{1/3} M\sqrt{2p} \leq 2.04 M\sqrt{p}$ for all $p \in \mathbb{Z}_{++}$. In this case, we have $\mathbb{E}\{|Z|^{2p}\} \leq (2.04 M\sqrt{2p})^{2p} = (2.04^2 M^2 2p)^p \leq (8.33 M^2 p)^p$. Also, by our assumption that $|\theta| \leq \frac{1}{7M}$, we get $8.33 e M^2 \theta^2 \leq \frac{8.33e}{49} \leq \frac{1}{2}$. Lastly, we can verify the inequality $e^{2x} \geq \frac{1}{1-x}$ for all $x \in [0, \frac{1}{2}]$. Thus, using the Taylor series expansion of e^x , we have

$$\begin{aligned} \mathbb{E}\{e^{\theta^2 Z^2}\} &= 1 + \sum_{p=1}^\infty \frac{\theta^{2p} \mathbb{E}\{|Z|^{2p}\}}{p!} \leq 1 + \sum_{p=1}^\infty \frac{(\theta^2 8.33 M^2 p)^p}{p!} \stackrel{(e)}{\leq} 1 + \sum_{p=1}^\infty \frac{(\theta^2 8.33 M^2 p)^p}{(p/e)^p} \\ &= 1 + \sum_{p=1}^\infty (\theta^2 8.33 e M^2)^p \stackrel{(f)}{=} \frac{1}{1 - \theta^2 8.33 e M^2} \stackrel{(g)}{\leq} e^{(7M\theta)^2}, \end{aligned}$$

where (e) holds because $p! \geq (p/e)^p$ by the Sterling approximation, (f) follows by $8.33 e M^2 \theta^2 \leq \frac{1}{2}$ and (g) holds because $8.33 e M^2 \theta^2 \leq \frac{1}{2}$, so $\frac{1}{1 - \theta^2 8.33 e M^2} \leq e^{2\theta^2 8.33 e M^2} \leq e^{49\theta^2 M^2}$. \blacksquare

In the proof, we round the fractions up to two decimals to strike a balance between clarity and tightness. We can round the fractions up to integers, in which case, the bounds would be looser.

Lemma L.6 *If Z is a mean-zero random variable that satisfies $\mathbb{E}\{e^{\lambda^2 Z^2}\} \leq e^{M^2\lambda^2}$ for all $|\lambda| \leq \frac{1}{M}$, then we have $\mathbb{E}\{e^{\theta Z}\} \leq e^{M^2\theta^2}$ for all $\theta \in \mathbb{R}$.*

Proof: Consider any $\theta \in \mathbb{R}$. We claim that $\mathbb{E}\{e^{\frac{\theta}{M}Z}\} \leq e^{\theta^2}$. First, assume that $|\theta| \leq 1$. Because $|\frac{\theta}{M}| \leq \frac{1}{M}$, using the assumption of the lemma with $\lambda = \frac{\theta}{M}$, we get $\mathbb{E}\{e^{\frac{\theta^2}{M^2}Z^2}\} \leq e^{\theta^2}$. For any $x \in \mathbb{R}$,

we have the inequality $e^x \leq x + e^{x^2}$, in which case, noting that $\mathbb{E}\{Z\} = 0$, we obtain $\mathbb{E}\{e^{\frac{\theta}{M}Z}\} \leq \mathbb{E}\{\frac{\theta}{M}Z\} + \mathbb{E}\{e^{\frac{\theta^2}{M^2}Z^2}\} \leq e^{\theta^2}$, establishing the claim when $|\theta| \leq 1$. Second, assume that $|\theta| \geq 1$. Because $\frac{1}{\sqrt{2M}} \leq \frac{1}{M}$, using the assumption of the lemma with $\lambda = \frac{1}{\sqrt{2M}}$, we get $\mathbb{E}\{e^{\frac{1}{2M^2}Z^2}\} \leq e^{\frac{1}{2}}$. We have the standard inequality $\frac{1}{2}a^2 + \frac{1}{2}b \geq ab$, which implies that $\frac{\theta^2}{2} + \frac{x^2}{2M^2} \geq \frac{\theta x}{M}$ for all $x \in \mathbb{R}$. In this case, we obtain the chain of inequalities $\mathbb{E}\{e^{\frac{\theta}{M}Z}\} \leq e^{\frac{\theta^2}{2}} \mathbb{E}\{e^{\frac{1}{2M^2}Z^2}\} \leq e^{\frac{\theta^2}{2}} e^{\frac{1}{2}} \leq e^{\theta^2}$, where the last inequality holds because we have $|\theta| \geq 1$. Therefore, the claim holds when $|\theta| \geq 1$ as well.

By the claim established in the previous paragraph, we have $\mathbb{E}\{e^{\frac{\theta}{M}Z}\} \leq e^{\theta^2}$ for all $\theta \in \mathbb{R}$, which is equivalent to having $\mathbb{E}\{e^{\theta Z}\} \leq e^{M^2\theta^2}$ for all $\theta \in \mathbb{R}$. ■

L.3 Simplifying the Fluid Approximation Under Independent Demands

Problem (2) has a simplified form under independent demands. Using the decision variables $\mathbf{x} = (x_{jt}^k : j \in \mathcal{J}, t \in \mathcal{T}, k \in \mathcal{K})$, we consider the problem

$$\begin{aligned} \max_{\mathbf{x} \in \mathbb{R}_+^{|\mathcal{J}|T\mathcal{K}}} \quad & \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j \mathbb{P}\{D^k \geq t\} x_{jt}^k \\ \text{st} \quad & \sum_{\ell=1}^{k-1} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq t\} x_{jt}^\ell + \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} x_{jt}^k \leq c_i \quad \forall i \in \mathcal{L}, k \in \mathcal{K} \\ & x_{jt}^k \leq \lambda_{jt}^k \quad \forall j \in \mathcal{J}, t \in \mathcal{T}, k \in \mathcal{K}. \end{aligned} \tag{29}$$

We shortly show that if the demands are independent, then problem (2) is equivalent to problem (29). Furthermore, problem (29) is useful to see that naive fluid approximations are inadequate to deal with multiple stages even under independent demands. In particular, the first constraint above uses the distribution of the demand for stages $1, \dots, k-1$, but it does not use the distribution of the demand for stage k . The specific form of this constraint is critical. A naive approach may use a constraint of the form $\sum_{\ell=1}^k \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq t\} x_{jt}^\ell \leq c_i$ or $\sum_{\ell=1}^k \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} x_{jt}^\ell \leq c_i$. We also shortly demonstrate that the former option does not yield an asymptotically tight fluid approximation, whereas the latter option does not even yield an upper bound on the optimal total expected revenue. Bai et al. (2023) give fluid approximations under random number of customer arrivals that occur in one stage. Because they have one stage, their fluid approximations do not capture dependence between the demands in different stages. Even under independent demands, using one of the two options earlier in this paragraph to extend their work to multiple stages yields fluid approximations that are not asymptotically tight.

Equivalence of Problems (2) and (29):

We show that if the demands in different stages are independent, then the optimal objective values of (2) and (29) are equal to each other and we can construct an optimal solution to one

problem by using an optimal solution to the other. First, letting $\bar{\mathbf{x}} = (\bar{x}_{jt}^k(q) : j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$ be an optimal solution to (2), we define the solution $\bar{\mathbf{y}} = (\bar{y}_{jt}^k : j \in \mathcal{J}, t \in \mathcal{T}, k \in \mathcal{K})$ to (29) as $\bar{y}_{jt}^k = \sum_{q \in \mathcal{T}} \mathbb{P}\{D^{k-1} = q\} \bar{x}_{jt}^k(q)$. We claim that the solution $\bar{\mathbf{y}}$ is feasible to problem (29) and provides an objective value that is equal to the optimal objective value of problem (2). For problem (29), the solution $\bar{\mathbf{y}}$ provides an objective value of $\sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j \mathbb{P}\{D^k \geq t\} \bar{y}_{jt}^k = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j \mathbb{P}\{D^k \geq t\} \sum_{q \in \mathcal{T}} \mathbb{P}\{D^{k-1} = q\} \bar{x}_{jt}^k(q)$. Because the demands are independent, we have $\mathbb{P}\{D^k \geq t\} \mathbb{P}\{D^{k-1} = q\} = \mathbb{P}\{D^k \geq t, D^{k-1} = q\}$, in which case, the last sum corresponds to the optimal objective value of problem (2). Thus, the solution $\bar{\mathbf{y}}$ provides an objective value for (29) that is equal to the optimal objective value of (2). Because $\bar{\mathbf{x}}$ is an optimal solution to problem (2), it satisfies the second constraint in this problem. Multiplying this constraint with $\mathbb{P}\{D^{k-1} = q\}$ and adding over all $q \in \mathcal{T}$, we obtain $\bar{y}_{jt}^k = \sum_{q \in \mathcal{T}} \mathbb{P}\{D^{k-1} = q\} \bar{x}_{jt}^k(q) \leq \lambda_{jt}^k$, in which case, it follows that the solution $\bar{\mathbf{y}}$ satisfies the second constraint in problem (29). Under independent demands, for $\ell = 1, \dots, k-1$, we have the identity $\sum_{q \in \mathcal{T}} \mathbb{P}\{D^\ell \geq s \mid D^k \geq t, D^{k-1} = q\} \mathbb{P}\{D^{k-1} = q\} = \mathbb{P}\{D^\ell \geq s\}$. In particular, if $\ell \leq k-2$, then this identity holds because D^ℓ, D^{k-1} and D^k are independent. If, on the other hand, $\ell = k-1$, then we have $\sum_{q \in \mathcal{T}} \mathbb{P}\{D^{k-1} \geq s \mid D^k \geq t, D^{k-1} = q\} \mathbb{P}\{D^{k-1} = q\} = \sum_{q \in \mathcal{T}} \mathbb{P}\{D^{k-1} \geq s, D^{k-1} = q\}$, where we use the fact that D^k and D^{k-1} are independent. The identity follows by noting that the last sum is equal to $\mathbb{P}\{D^{k-1} \geq s\}$. Because $\bar{\mathbf{x}}$ is an optimal solution to problem (2), it satisfies the first constraint in this problem with $t = T$. Multiplying this constraint with $\mathbb{P}\{D^{k-1} = q\}$ and adding over all $q \in \mathcal{T}$, as well as using the fact that the demands are independent, we obtain the inequality

$$\begin{aligned} \sum_{q \in \mathcal{T}} \sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^{\ell-1} = p\} \mathbb{P}\{D^\ell \geq s \mid D^k \geq t, D^{k-1} = q\} \mathbb{P}\{D^{k-1} = q\} \bar{x}_{js}^\ell(p) \\ + \sum_{q \in \mathcal{T}} \sum_{s \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^{k-1} = q\} \bar{x}_{js}^k(q) \leq c_i. \end{aligned}$$

Arranging the terms on the left side of the inequality above, as well as noting the fact that we have $\bar{y}_{jt}^k = \sum_{q \in \mathcal{T}} \mathbb{P}\{D^{k-1} = q\} \bar{x}_{jt}^k(q)$ by definition of \bar{y}_{jt}^k and the identity established just before the inequality above, the left side of the inequality above is $\sum_{\ell=1}^{k-1} \sum_{s \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq s\} \bar{y}_{js}^\ell + \sum_{s \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \bar{y}_{js}^k$. Thus, the solution $\bar{\mathbf{y}}$ satisfies the first constraint in problem (29) as well. In this case, the solution $\bar{\mathbf{y}}$ is feasible to problem (29) and the desired claim follows.

Second, letting $\bar{\mathbf{y}} = (\bar{y}_{jt}^k : j \in \mathcal{J}, t \in \mathcal{T}, k \in \mathcal{K})$ be an optimal solution to problem (29), we define the solution $\bar{\mathbf{x}} = (\bar{x}_{jt}^k(q) : j \in \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$ to problem (2) as $\bar{x}_{jt}^k(q) = \bar{y}_{jt}^k$. In this case, we can follow an argument analogous to, but slightly simpler than, the one in the previous paragraph to show that $\bar{\mathbf{x}}$ is a feasible solution to problem (2) and provides an objective value that is equal to the optimal objective value of problem (29). Therefore, by the discussion so far in this section, if the

demands in different stages are independent, then the optimal objective values of problems (2) and (29) are equal to each other and we can construct an optimal solution to one problem by using an optimal solution to the other. In the rest of this section, we give two problem instances. We use the first problem instance to demonstrate that if we replace the first constraint in problem (29) with $\sum_{\ell=1}^k \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq t\} x_{jt}^\ell \leq c_i$, then the fluid approximation in (29) is not asymptotically tight. We use the second problem instance to demonstrate that if we replace the first constraint in problem (29) with $\sum_{\ell=1}^k \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} x_{jt}^\ell \leq c_i$, then the optimal objective value of problem (29) does not provide an upper bound on the optimal total expected revenue.

Losing the Asymptotic Tightness of the Fluid Approximation:

We give a problem instance with independent demands to demonstrate that if we replace the first constraint in (29) with $\sum_{\ell=1}^k \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq t\} x_{jt}^\ell \leq c_i$, then the fluid approximation in (29) is not asymptotically tight. In our problem instance, there is one stage, one resource and one product. Thus, we have $K = 1$, $\mathcal{L} = \{1\}$ and $\mathcal{J} = \{1\}$. The capacity of the resource is $c_1 = \alpha^3$ for some $\alpha \in \mathbb{Z}_+$. The revenue of the product is $f_1 = 1$. The demand in the single stage has two possible values with $\mathbb{P}\{D^1 = \alpha\} = 1 - \frac{1}{\alpha^2}$ and $\mathbb{P}\{D^1 = \alpha^4\} = \frac{1}{\alpha^2}$. At each time period, there is a request for the product with probability one. We have $K = 1$ and $c_{\min} = \alpha^3$. Because there is one resource, we have $L = 1$. As a function of the parameter α , using OPT^α to denote the optimal total expected revenue, APX^α to denote the total expected revenue of the approximate policy and $\bar{Z}_{\text{LP}}^\alpha$ to denote the optimal objective value of (29), by Theorem L.1, we have $\frac{\text{OPT}^\alpha}{\bar{Z}_{\text{LP}}^\alpha} \geq \frac{\text{APX}^\alpha}{\bar{Z}_{\text{LP}}^\alpha} \geq 1 - 4 \frac{\sqrt{\alpha^3 \log \alpha^3}}{\alpha^3} - \frac{1}{\alpha^3}$. By Theorem 3.1, we have $1 \geq \frac{\text{OPT}^\alpha}{\bar{Z}_{\text{LP}}^\alpha}$. Thus, as α becomes arbitrarily large, the relative gap between the optimal total expected revenue and the optimal objective value of problem (29) vanishes.

We proceed to demonstrating that if we replace the first constraint in problem (29) with $\sum_{\ell=1}^k \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq t\} x_{jt}^\ell \leq c_i$, then the optimal total expected revenue and the optimal objective value of problem (29) satisfy $\frac{\text{OPT}^\alpha}{\bar{Z}_{\text{LP}}^\alpha} = \frac{2\alpha^2 - 1}{\alpha^3 + \alpha^2 - 1}$. Therefore, as α becomes arbitrarily large, the relative gap between the optimal total expected revenue and the optimal objective value of problem (29) becomes arbitrarily large, so we lose the asymptotic tightness of the fluid approximation. We compute the optimal total expected revenue. Because there is a single product, it is optimal to accept the requests for the product as much as the capacity allows. Noting that there is a product request at each time period, if $D^1 = \alpha$, then the capacity of α^3 allows us to accept all product requests at α time periods and we obtain a total expected revenue of α . If $D^1 = \alpha^4$, then we can sell all of the capacity by using the product requests at α^4 time periods, so noting that the capacity of the resource is α^3 , we obtain a total expected revenue of α^3 . Thus, the optimal total expected revenue is $\text{OPT}^\alpha = \mathbb{P}\{D^1 = \alpha\} \alpha + \mathbb{P}\{D^1 = \alpha^4\} \alpha^3 = (1 - \frac{1}{\alpha^2}) \alpha +$

$\frac{1}{\alpha^2} \alpha^3 = \frac{2\alpha^2-1}{\alpha}$. For our problem instance, if we replace the first constraint in problem (29) with $\sum_{\ell=1}^k \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq t\} x_{jt}^\ell \leq c_i$, then this problem is given by

$$\begin{aligned} \bar{Z}_{\text{LP}}^\alpha &= \max \quad \sum_{t=1}^{\alpha} x_{1t}^1 + \sum_{t=\alpha+1}^{\alpha^4} \frac{1}{\alpha^2} x_{1t}^1 \\ \text{st} \quad &\sum_{t=1}^{\alpha} x_{1t}^1 + \sum_{t=\alpha+1}^{\alpha^4} \frac{1}{\alpha^2} x_{1t}^1 \leq \alpha^3 \\ &x_{1t}^1 \in [0, 1] \quad \forall t = 1, \dots, \alpha^4. \end{aligned}$$

Because $\alpha + (\alpha^4 - \alpha) \frac{1}{\alpha^2} \leq \alpha^3$ for all $\alpha \in \mathbb{Z}_+$, setting all of the decision variables to one yields a feasible, as well as an optimal, solution, so $\bar{Z}_{\text{LP}}^\alpha = \alpha + (\alpha^4 - \alpha) \frac{1}{\alpha^2} = \frac{\alpha^3 + \alpha^2 - 1}{\alpha}$.

By the discussion in the previous paragraph, we have $\text{OPT}^\alpha = \frac{2\alpha^2-1}{\alpha}$ and $\bar{Z}_{\text{LP}}^\alpha = \frac{\alpha^3 + \alpha^2 - 1}{\alpha}$ for our problem instance, which implies that $\frac{\text{OPT}^\alpha}{\bar{Z}_{\text{LP}}^\alpha} = \frac{2\alpha^2-1}{\alpha^3 + \alpha^2 - 1}$.

Losing the Upper Bound from the Fluid Approximation:

We give a problem instance with independent demands to demonstrate that if we replace the first constraint in (29) with $\sum_{\ell=1}^k \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} x_{jt}^\ell \leq c_i$, then the optimal objective value of problem (29) is not an upper bound on the optimal total expected revenue. In our problem instance, there are two stages, one product and one resource, so $K = 2$, $\mathcal{L} = \{1\}$ and $\mathcal{J} = \{1\}$. The capacity of the resource is 3. The revenue of the product is $f_1 = 1$. The demand in each of the two stages has two possible values with $\mathbb{P}\{D^1 = 1\} = \mathbb{P}\{D^1 = 2\} = \mathbb{P}\{D^2 = 1\} = \mathbb{P}\{D^2 = 2\} = \frac{1}{2}$. At each time period in each stage, we have a request for the product with probability one. We proceed to computing the optimal total expected revenue. Because there is a single product, it is optimal to accept the requests for the product as much as the capacity allows. With probability $\frac{1}{2} \times \frac{1}{2} = \frac{1}{4}$, the total number of requests for the product over the selling horizon is two, in which case, we make a revenue of two. With probability $1 - \frac{1}{4} = \frac{3}{4}$, the total number of requests for the product over the selling horizon is three or more, in which case, noting that the capacity of the resource is three, we make a revenue of three. Therefore, the optimal total expected revenue is given by $\text{OPT} = \frac{1}{4} \times 2 + \frac{3}{4} \times 3 = \frac{11}{4}$. For our problem instance, if we replace the first constraint in problem (29) with $\sum_{\ell=1}^k \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} x_{jt}^\ell \leq c_i$, then this problem is given by

$$\begin{aligned} \bar{Z}_{\text{LP}} &= \max \quad x_{11}^1 + \frac{1}{2} x_{12}^1 + x_{11}^2 + \frac{1}{2} x_{12}^2 \\ \text{st} \quad &x_{11}^1 + x_{12}^1 \leq 3 \\ &x_{11}^1 + x_{12}^1 + x_{11}^2 + x_{12}^2 \leq 3 \\ &x_{11}^1, x_{12}^1, x_{11}^2, x_{12}^2 \in [0, 1]. \end{aligned}$$

Setting all of the decision variables other than x_{12}^2 to one yields an optimal solution to the problem above, so $\bar{Z}_{\text{LP}} = 1 + \frac{1}{2} + 1 = \frac{5}{2}$. Therefore, for our problem instance, we have $\bar{Z}_{\text{LP}} = \frac{5}{2} < \frac{11}{4} = \text{OPT}$.

Appendix M: Extension to Assortment Offering and Pricing Decisions

We give the extension of our results to the case where we make assortment offering or pricing decisions. For economy of space, we only point out the adjustments needed in our development.

Problem Formulation:

We give the pieces of notation that we need in addition to those given in Section 2. We pick an assortment of products to offer to the customer arriving at each time period. The customer either makes a choice within the offered assortment or leaves without making a purchase. Our goal is to find a policy to decide which assortment to offer to each customer so that we maximize the total expected revenue over the selling horizon. Given that we offer the assortment of products $S \subseteq \mathcal{J}$, we use $\phi_{jt}^k(S)$ to denote the probability that the customer arriving at time period t in stage k chooses product j . We have $\phi_{jt}^k(S) = 0$ for all $j \in \mathcal{J} \setminus S$. With probability $1 - \sum_{j \in \mathcal{J}} \phi_{jt}^k(S)$, the customer leaves without making a purchase. As a function of the remaining capacities $\mathbf{y} = (y_i : i \in \mathcal{L}) \in \mathbb{Z}_+^{|\mathcal{L}|}$, the set of feasible decisions is $\mathcal{F}(\mathbf{y}) = \{S \subseteq \mathcal{J} : a_{ij} \mathbf{1}(j \in S) \leq y_i \forall i \in \mathcal{L}, j \in \mathcal{J}\}$, ensuring that if we want to offer an assortment that includes product j and the product uses resource i , then we need to have at least one unit of remaining capacity for resource i . We can find the optimal policy by computing the value functions $\{J_t^k : t \in \mathcal{T}, k \in \mathcal{K}\}$ through the dynamic program

$$J_t^k(\mathbf{y}, q) = \max_{S \in \mathcal{F}(\mathbf{y})} \left\{ \sum_{j \in \mathcal{J}} \phi_{jt}^k(S) \left\{ f_j + \theta_t^k(q) J_{t+1}^k(\mathbf{y} - \mathbf{a}_j, q) + (1 - \theta_t^k(q)) J_1^{k+1}(\mathbf{y} - \mathbf{a}_j, t) \right\} \right. \\ \left. + \left(1 - \sum_{j \in \mathcal{J}} \phi_{jt}^k(S) \right) \left\{ \theta_t^k(q) J_{t+1}^k(\mathbf{y}, q) + (1 - \theta_t^k(q)) J_1^{k+1}(\mathbf{y}, t) \right\} \right\},$$

with the boundary condition that $J_1^{K+1} = 0$. The demand in the stage just before the beginning of the selling horizon is fixed at \widehat{D}^0 , so the optimal total expected revenue is $\text{OPT} = J_1^1(\mathbf{c}, \widehat{D}^0)$.

Our formulation with the assortment offering decisions in the dynamic program above can capture pricing decisions as a special case. To use the formulation above to capture pricing decisions, we create multiple copies of each product, where a copy of a product corresponds to offering the product at a particular price. We refer to each copy of a product as a virtual product. If the set of possible price levels for a product is \mathcal{P} , then we have $|\mathcal{J}| \times |\mathcal{P}|$ virtual products and we index the virtual products by $\{(j, p) \in \mathcal{J} \times \mathcal{P}\}$. In the pricing problem, we pick a subset of virtual products to offer to a customer with the constraint that if we offer the virtual product (j, p) , then we cannot offer another virtual product of the form (j, q) with $q \in \mathcal{P} \setminus \{p\}$. This constraint ensures that if we offer a product, then we offer it at only one price level. Given that we offer the assortment of virtual products $S \subseteq \mathcal{J} \times \mathcal{P}$, we use $\phi_{jpt}^k(S)$ to denote probability that the customer arriving at time period t in stage k chooses virtual product (j, p) . In this case, we can continue using the model in the previous paragraph with the understanding that virtual products play the role of

products. Thus, if there is a set of discrete possible price levels that we can charge for each product, then our discussion in this section applies to the pricing setting.

Fluid Approximation:

We use the decision variable $x_t^k(S, q)$ to capture the probability of offering assortment S at time period t in stage k given that the demand in the previous stage was q . Consider the problem

$$\begin{aligned} \bar{Z}_{\text{LP}} = \max_{\mathbf{x} \in \mathbb{R}_+^{2^{|\mathcal{J}|} T^2 K}} & \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{S \subseteq \mathcal{J}} \sum_{q \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j \mathbb{P}\{D^k \geq t, D^{k-1} = q\} \phi_{jt}^k(S) x_t^k(S, q) \\ \text{st} & \sum_{\ell=1}^{k-1} \sum_{r \in \mathcal{T}} \sum_{S \subseteq \mathcal{J}} \sum_{p \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^\ell \geq r, D^{\ell-1} = p \mid D^k \geq t, D^{k-1} = q\} \phi_{jr}^\ell(S) x_r^\ell(S, p) \\ & + \sum_{r=1}^t \sum_{S \subseteq \mathcal{J}} \sum_{j \in \mathcal{J}} a_{ij} \phi_{jr}^k(S) x_r^k(S, q) \leq c_i \quad \forall i \in \mathcal{L}, t, q \in \mathcal{T}, k \in \mathcal{K} \\ & \sum_{S \subseteq \mathcal{J}} x_t^k(S, q) = 1 \quad \forall t, q \in \mathcal{T}, k \in \mathcal{K}, \end{aligned} \quad (30)$$

where we use the vector of decision variables $\mathbf{x} = (x_t^k(S, q) : S \subseteq \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$. Problem (30) is the analogue of problem (2) under assortment offering decisions. The left side of the first constraint corresponds to the total expected capacity consumption of resource i up to and including time period t in stage k conditional on the fact that the demand in stage k is at least t and the demand in stage $k-1$ is q . The second constraint ensures we offer some assortment with probability one at each time period in each stage, but we allow offering an empty assortment. We can show that the optimal objective value of the problem above is an upper bound on the optimal total expected revenue, so $\bar{Z}_{\text{LP}} \geq \text{OPT}$. This result is the analogue of Theorem 3.1. We can show this result using any of the three approaches discussed at the end of Section 3, which include Lagrangian relaxation, linear value function approximations and sample paths of the optimal policy.

Approximate Policy and Asymptotic Optimality:

We can use problem (30) to construct an approximate policy. In particular, we solve problem (30) once at the beginning of the selling horizon. Letting $\bar{\mathbf{x}} = (\bar{x}_t^k(S, q) : S \subseteq \mathcal{J}, t, q \in \mathcal{T}, k \in \mathcal{K})$ be an optimal solution to this problem, the approximate policy makes its decisions as follows. Using $\gamma \in [0, 1]$ to denote a tuning parameter, if the demand in stage $k-1$ was q , then we offer assortment S at time period t in stage k with probability $\gamma \bar{x}_t^k(S, q)$. With probability $1 - \gamma \sum_{S \subseteq \mathcal{J}} \bar{x}_t^k(S, q)$, we offer the empty assortment. If the customer chooses a product within the offered assortment for which we do not have remaining resource capacity to serve, then the customer leaves without a purchase. Otherwise, the customer purchases the chosen product. We use $c_{\min} = \min_{i \in \mathcal{L}} c_i$ to denote the smallest resource capacity and $L = \max_{j \in \mathcal{J}} \sum_{i \in \mathcal{L}} a_{ij}$ to denote the maximum number of resources used by a product. In this case, letting APX be the total expected revenue of the

approximate policy, recalling that OPT is the optimal total expected revenue and \bar{Z}_{LP} is the optimal objective value of the linear program in (30), we can show that

$$\frac{\text{APX}}{\text{OPT}} \geq \frac{\text{APX}}{\bar{Z}_{\text{LP}}} \geq \max \left\{ \frac{1}{4L}, \left(1 - 4 \frac{\sqrt{(c_{\min} + \frac{1}{\epsilon^6} (K-1)) \log c_{\min}}}{c_{\min}} - \frac{L}{c_{\min}} \right) \right\}.$$

This result is the analogue of Theorem 4.1 under assortment offering decisions. The interpretation of the performance guarantee above is identical to that of the performance guarantee in Theorem 4.1. In many network revenue management settings, the numbers of resources and products can be large, but the number of resources used by a particular product is small. Therefore, the first part of the performance guarantee ensures that the approximate policy has a constant-factor performance guarantee when L is uniformly bounded. On the other hand, considering a regime where we scale both the number of stages and resource capacities with the same rate θ , the second part of the performance guarantee ensures that the relative gap between the total expected revenue of the approximate policy and the optimal total expected revenue converges to one with rate $1 - \frac{1}{\sqrt{\theta}}$. We proceed to giving an outline for the proof of the performance guarantee for our approximate policy under assortment offering decisions.

Performance Guarantee:

The performance guarantee under assortment offering decisions follows from the same outline in Section 5. For each $k \in \mathcal{K}$ and $t \in \mathcal{T}$, we define four classes of Bernoulli random variables.

- *Demand in Each Stage.* For each $q \in \mathcal{T}$, the random variable $\Psi_t^k(q)$ takes value one if we reach time period t in stage k before this stage is over and the demand in stage $k-1$ is q . In other words, we have $\Psi_t^k(q) = \mathbf{1}(D^k \geq t, D^{k-1} = q)$.

- *Policy Decision.* For each $S \subseteq \mathcal{J}$ and $q \in \mathcal{T}$, the random variable $X_t^k(S, q)$ takes value one if the approximate policy offers assortment S at time period t in stage k when the demand in stage $k-1$ was q . By our approximate policy, $\mathbb{P}\{X_t^k(S, q) = 1\} = \gamma \bar{x}_t^k(S, q)$.

- *Customer Choice.* For each $S \subseteq \mathcal{J}$ and $j \in \mathcal{J}$, given that the approximate policy offers the assortment S , the random variable $\Phi_{jt}^k(S)$ takes value one if the customer arriving at time period t in stage k purchases product j . We have $\mathbb{P}\{\Phi_{jt}^k(S) = 1\} = \phi_{jt}^k(S)$.

- *Availability.* For each $j \in \mathcal{J}$, the random variable G_{jt}^k takes value one if we have enough capacity to accept a request for product j at time period t in stage k under the approximate policy. Instead of calculating the probability $\mathbb{P}\{G_{jt}^k = 1\}$, we lower bound $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\}$.

The equality in (3) continues to hold as long we replace $\bar{x}_{jt}^k(q)$ in this equality with $\sum_{S \subseteq \mathcal{J}} \phi_{jt}^k(S) \bar{x}_t^k(S, q)$. Thus, if we can show that $\mathbb{P}\{G_{jt}^k = 1 \mid \Psi_t^k(q) = 1\} \geq \alpha$, then $\text{APX} \geq \gamma \alpha \bar{Z}_{\text{LP}}$.

Under the approximate policy, the sales for product j at time period t in stage k is given by $\sum_{q \in \mathcal{T}} \sum_{S \subseteq \mathcal{J}} \Psi_t^k(q) G_{jt}^k \Phi_{jt}^k(S) X_t^k(S, q)$, where we use the fact that there is a sale for product j at time period t in stage k if and only if we reach time period t in stage k , we have enough capacity to accept a request for product j and the customer chooses product j out of the offered assortment. Thus, we can use $\sum_{q \in \mathcal{T}} \sum_{S \subseteq \mathcal{J}} \Psi_t^k(q) \Phi_{jt}^k(S) X_t^k(S, q)$ as an upper bound on the sales for product j at time period t in stage k . In this case, $\sum_{j \in \mathcal{J}} \sum_{q \in \mathcal{T}} \sum_{S \subseteq \mathcal{J}} a_{ij} \Psi_t^k(q) \Phi_{jt}^k(S) X_t^k(S, q)$ is an upper bound on the capacity consumption of resource i at time period t in stage k . Letting $N_{it}^k(q) = \sum_{S \subseteq \mathcal{J}} \sum_{j \in \mathcal{J}} a_{ij} \Phi_{jt}^k(S) X_t^k(S, q)$, we express the last upper bound as $\sum_{q \in \mathcal{T}} \Psi_t^k(q) N_{it}^k(q)$, in which case, the chain of inequalities in (4) continues to hold. Lemmas 5.1, 5.2 and 5.3 go through with no modifications when we use the definition of $N_{it}^k(q)$ given in this paragraph. We obtain the performance guarantee for our approximate policy under assortment offering decisions by following precisely the same outline in the proof of Theorem 4.1. To sum up, if we decide whether to accept or reject each product request, then we use $\sum_{q \in \mathcal{T}} \Psi_t^k(q) A_{jt}^k X_{jt}^k$ as an upper bound on the sales for product j at time period t in stage k . If we make assortment offering decisions, then we use $\sum_{q \in \mathcal{T}} \sum_{S \subseteq \mathcal{J}} \Psi_t^k(q) \Phi_{jt}^k(S) X_t^k(S, q)$ as an upper bound on the sales for product j at time period t in stage k . Once we draw this parallel, the performance guarantee under assortment offering decisions follows precisely the same outline in the proof of Theorem 4.1.

At the beginning of this section, we explain that our formulation with assortment offering decision captures pricing decisions as a special case. To elaborate further, if the set of possible price levels for a product is \mathcal{P} , then the set of virtual products is $\mathcal{J} \times \mathcal{P}$. Offering virtual product (j, p) implies that we offer product j at price level p . In the pricing problem, we choose an assortment of virtual product to offer at each time period in each stage and the arriving customers make a choice among the offered virtual products. A product, if offered, can be offered at a single price. Therefore, if the remaining capacities of the resources at any time period in any stage are given by the vector $\mathbf{y} = (y_i : i \in \mathcal{L})$, then the set of feasible assortments of virtual products that we can offer corresponds to $\mathcal{F}(\mathbf{y}) = \{S \subseteq \mathcal{J} \times \mathcal{P} : a_{ij} \mathbf{1}(S \cap (\{j\} \times \mathcal{P})) \leq y_i \forall i \in \mathcal{L}, j \in \mathcal{J}, |S \cap (\{j\} \times \mathcal{P})| \leq 1 \forall j \in \mathcal{J}\}$, where the first constraint ensures that if we want to offer an assortment that includes a virtual product corresponding to product j and product j uses resource i , then we need to have at least one unit of remaining capacity for resource i , whereas the second constraint ensures that we can offer at most one virtual product corresponding to product j . To capture the reactions of the customers to prices, given that we offer the assortment of virtual products $S \subseteq \mathcal{J} \times \mathcal{P}$, we use $\phi_{jpt}^k(S)$ to denote the probability that the customer arriving at time period t in stage k chooses virtual product (j, p) . In this case, our model under assortment offering decisions applies to pricing decisions, as long as we view each virtual product as a product under assortment offering decisions.

Appendix N: Computational Experiments

We conduct a numerical study to investigate the benefits from our fluid approximation when the demands in different stages have arbitrary distributions and are dependent on each other.

Experimental Setup: In our computational experiments, we consider an airline network with one hub and three spokes. We have a flight leg that connects the hub to each spoke and each spoke to the hub. Thus, there are six flight legs. We have a high-fare and a low-fare itinerary that connects each origin-destination pair. Thus, there are $2 \times 4 \times 3 = 24$ itineraries. In this setting, the resources correspond to the flight legs, whereas the products correspond to the itineraries. The itineraries between the hub and a spoke are direct, each using a single resource, whereas the itineraries between two spokes connect at the hub, each using two resources. To generate the revenues associated with the itineraries, we place the hub at the center of a 100×100 square. We place the spokes over the square uniformly at random. The revenue of a low-fare itinerary is the Euclidean distance between the origin and destination locations of the itinerary. The revenue of a high-fare itinerary is κ times the revenue of the corresponding low-fare itinerary. The preceding discussion explains our approach for setting up the resources and products, along with the set of resources used by each product. We proceed to explaining our approach for generating the distribution of the demand random variables in each stage and the correlation structure between the demand random variables.

We have K stages in the selling horizon. We will vary the parameter K . The demand in each stage has a truncated and discretized log-normal distribution. In particular, given that the demand in stage $k - 1$ takes the value \bar{D}^{k-1} , the mean and standard deviation of the demand in stage k are, respectively, $\mu^k(\bar{D}^{k-1}) = \rho \bar{D}^{k-1} + (1 - \rho) 100$ and $\sigma^k(\bar{D}^{k-1}) = 0.3(\rho \bar{D}^{k-1} + (1 - \rho) 100)$. In this case, given that the demand in stage $k - 1$ takes the value \bar{D}^{k-1} , the demand in stage k is obtained by rounding a log-normal random variable with mean $\mu^k(\bar{D}^{k-1})$ and standard deviation $\sigma^k(\bar{D}^{k-1})$ up to the closest integer and truncating at $\lceil \mu^k + 3\sigma^k \rceil$. We fix D^0 at $\hat{D}^0 = 100$. Therefore, the demand in each stage, before rounding and truncating, always has the mean of 100 and coefficient of variation of 0.3. The parameter ρ controls the correlation between the demand random variables in successive stages. Because a linear transformation of a log-normal random variable is not a log-normal random variable, the correlation coefficient between the demands in successive stages may not be necessarily ρ . After generating each test problem, however, we computed the correlation coefficient between the demands in successive stages and the correlation coefficient came out to be quite close to ρ . We will vary the parameter ρ .

To come up with the request arrival probabilities $\{\lambda_{jt}^k : j \in \mathcal{J}, t \in \mathcal{T}, k \in \mathcal{K}\}$, we assume that the requests for the low-fare itineraries tend to arrive towards the beginning of the selling horizon,

whereas the requests for the high-fare itineraries tend to arrive towards the end of the selling horizon. In this way, we generate test problems in which it is important to protect capacities for the high-fare itinerary requests that tend to arrive later. In particular, the probability of getting a request for a low-fare itinerary linearly decreases over time. There are K stages, each with at most T time periods. We sample a threshold from the uniform distribution over $\{\lceil \frac{1}{2}KT \rceil, \dots, \lceil \frac{2}{3}KT \rceil\}$ so that the requests for the high-fare itineraries arrive only after the threshold and the probability of getting a request for a high-fare itinerary linearly increases over time. We use the following approach to come up with such request arrival probabilities.

For each origin-destination pair (f, g) , we sample ζ_{fg} from the uniform distribution over $[0, 1]$. Using N to denote the set of all locations, we set $\gamma_{fg} = \frac{\zeta_{fg}}{\sum_{k \in N} \sum_{\ell \in N \setminus \{k\}} \zeta_{k\ell}}$ so that the parameters $\{\gamma_{fg} : f \in N, g \in N \setminus \{f\}\}$ are normalized to add up to one. At any time period in any stage, we have a request for an itinerary that connects origin-destination pair (f, g) with probability γ_{fg} . There is a high-fare and a low-fare itinerary that connects each origin-destination pair. The probability that we have a request for a low-fare itinerary linearly increases over time. The probability that we have a request for a high-fare itinerary is zero until the threshold and increases linearly over time after the threshold. Therefore, for each origin-destination pair (f, g) , we sample the threshold τ_{fg} from the uniform distribution over $\{\lceil \frac{1}{2}KT \rceil, \dots, \lceil \frac{2}{3}KT \rceil\}$. The function $F(t) = \frac{KT+1-t}{KT}$ is decreasing for $t \in [1, KT]$, whereas the function $H_{fg}(t) = \frac{[t-\tau_{fg}]^+}{KT-\tau_{fg}}$ takes the value zero for $t \in [1, \tau_{fg}]$ and is increasing for $t \in [\tau_{fg}, KT]$. In this case, noting that there are $(k-1)T$ time periods before stage k , if product j corresponds to a low-fare itinerary that connects origin-destination pair (f, g) , then we set $\lambda_{jt}^k = \gamma_{fg} \frac{F((k-1)T+t)}{F((k-1)T+t)+H_{fg}((k-1)T+t)}$, whereas if product j corresponds to a high-fare itinerary that connects origin-destination pair (f, g) , then we set $\lambda_{jt}^k = \gamma_{fg} \frac{H_{fg}((k-1)T+t)}{F((k-1)T+t)+H_{fg}((k-1)T+t)}$.

Once we generate the request arrival probabilities, the total expected demand for the capacity on flight leg i is $\xi_i = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} a_{ij} \mathbb{P}\{D^k \geq t\} \lambda_{jt}^k$. We set the capacity of flight leg i as $c_i = \lceil \xi_i / \beta \rceil$, where the parameter β controls the tightness of the capacities. Thus, the total expected demand for the capacity on a flight leg exceeds the capacity of the flight leg by a factor of β . We fix $\kappa = 8$ and $\beta = 1.6$. We experimented with different values for these two parameters and our results remained qualitatively the same. We build on the log-normal distribution to capture the demand random variables. Using this distribution, we can increase the coefficient of variation of the demand as much as we like and not be concerned about negative values.

We vary $K \in \{5, 10, 15, 20, 25, 30\}$ and $\rho \in \{0.2, 0.4, 0.6, 0.8\}$. For each parameter combination (K, ρ) , we generate a test problem using the approach described so far in this section.

Benchmark: We use a benchmark fluid approximation that uses only the expected values of the demands. The total expected demand for product j is $\sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \mathbb{P}\{D^k \geq t\} \lambda_{jt}^k$. Using the

decision variable w_j to capture the total number of requests for product j that we accept, using the vector $\mathbf{w} = (w_j : j \in \mathcal{J})$, we consider the linear program

$$\max_{\mathbf{w} \in \mathbb{R}_+^{|\mathcal{J}|}} \left\{ \sum_{j \in \mathcal{J}} f_j w_j : \sum_{j \in \mathcal{J}} a_{ij} w_j \leq c_i \quad \forall i \in \mathcal{L}, \quad w_j \leq \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \mathbb{P}\{D^k \geq t\} \lambda_{jt}^k \quad \forall j \in \mathcal{J} \right\}. \quad (31)$$

By the first constraint, the expected capacity consumption of each resource does not exceed its capacity. By the second constraint, the expected number of accepted requests for each product does not exceed its expected demand. The optimal objective value of (31) is an upper bound on the optimal total expected revenue. In particular, letting $\bar{\mathbf{x}}$ be an optimal solution to (2), we set $\bar{w}_j = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \sum_{q \in \mathcal{T}} \mathbb{P}\{D^k \geq t, D^{k-1} = q\} \bar{x}_{jt}^k(q)$. Considering the first constraint in (2) for $k = K$, if we multiply this constraint with $\mathbb{P}\{D^k = t, D^{k-1} = q\}$ and add over all $t, q \in \mathcal{T}$, then we can verify that $\bar{\mathbf{w}}$ satisfies the first constraint in (31). Considering the second constraint in (2), if we multiply this constraint with $\mathbb{P}\{D^k \geq t, D^{k-1} = q\}$ and add over all $t, q \in \mathcal{T}$, then we can verify that $\bar{\mathbf{w}}$ satisfies the second constraint in (31). In this way, the solution $\bar{\mathbf{w}}$ is feasible to problem (31) and provides an objective value that is equal to the optimal objective value of problem (2). Thus, the optimal objective value of (31) is an upper bound on that of (2), which is, in turn, an upper bound on the optimal total expected revenue. However, at the end of this section, we demonstrate that the upper bound from (31) is not asymptotically tight. Using (31), we can also obtain an approximate policy. Letting $\bar{\mathbf{w}}$ be an optimal solution to (31), the approximate policy is willing to accept a request for product j at any time period with probability $\bar{w}_j / \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \mathbb{P}\{D^k \geq t\} \lambda_{jt}^k$.

Computational Results: We refer to the fluid approximation in (2) as PRF to indicate that this fluid approximation uses information on the full probability distribution of the demands, whereas we refer to the fluid approximation in (31) as EXF to indicate that this fluid approximation uses information on only the expected values of the demands. In our computational experiments, we generate 24 test problems using the approach discussed earlier in this section. For each test problem, we solve PRF and EXF to compute the upper bounds on the optimal total expected revenue provided by the two fluid approximations. Furthermore, we simulate the decisions of the approximate policies provided by PRF and EXF for 1000 sample paths to estimate the total expected revenues of these two approximate policies. We use common random numbers when simulating the decisions of the two approximate policies. We give our computational results in Table 2. In this table, the first column gives the parameter combinations for our test problems using the pair (K, ρ) . The second column gives the upper bound on the optimal total expected revenue provided by PRF, whereas the third column gives the total expected revenue of the approximate policy from PRF. The fourth column gives the ratio between the total expected revenue of the approximate policy and the upper bound provided by PRF multiplied by 100. The fifth, sixth and seventh columns give

Params. (K, ρ)	PRF			EXF			Bnd. Gap	Plcy. Gap
	Bound	Policy	Ratio	Bound	Policy	Ratio		
(5, 0.2)	32,752	30,185	92.16%	34,533	28,206	81.68%	5.16%	7.01%
(5, 0.4)	31,755	28,700	90.38%	33,529	26,603	79.34%	5.29%	7.88%
(5, 0.6)	30,244	26,866	88.83%	31,983	24,404	76.30%	5.44%	10.09%
(5, 0.8)	27,690	24,285	87.70%	29,449	20,492	69.58%	5.97%	18.51%
(10, 0.2)	71,923	68,361	95.05%	74,333	65,324	87.88%	3.24%	4.65%
(10, 0.4)	69,342	64,991	93.73%	72,099	61,155	84.82%	3.82%	6.27%
(10, 0.6)	64,802	59,938	92.49%	67,861	55,599	81.93%	4.51%	7.80%
(10, 0.8)	55,250	49,893	90.30%	58,482	42,911	73.37%	5.53%	16.27%
(15, 0.2)	111,579	106,140	95.13%	114,223	102,259	89.53%	2.31%	3.80%
(15, 0.4)	107,571	102,204	95.01%	110,808	97,262	87.78%	2.92%	5.08%
(15, 0.6)	100,060	93,949	93.89%	104,080	87,636	84.20%	3.86%	7.20%
(15, 0.8)	82,458	76,242	92.46%	87,337	66,527	76.17%	5.59%	14.60%
(20, 0.2)	151,270	144,609	95.60%	154,021	140,783	91.40%	1.79%	2.72%
(20, 0.4)	145,995	139,229	95.37%	149,455	133,671	89.44%	2.31%	4.16%
(20, 0.6)	135,759	129,531	95.41%	140,308	122,021	86.97%	3.24%	6.15%
(20, 0.8)	110,234	105,274	95.50%	116,391	92,082	79.11%	5.29%	14.33%
(25, 0.2)	191,025	181,896	95.22%	193,835	177,921	91.79%	1.45%	2.23%
(25, 0.4)	184,421	176,305	95.60%	188,005	170,534	90.71%	1.91%	3.38%
(25, 0.6)	171,628	164,427	95.80%	176,483	156,205	88.51%	2.75%	5.26%
(25, 0.8)	138,672	131,287	94.67%	145,754	117,589	80.68%	4.86%	11.65%
(30, 0.2)	230,756	220,695	95.64%	233,608	216,844	92.82%	1.22%	1.78%
(30, 0.4)	222,995	214,150	96.03%	226,656	207,869	91.71%	1.62%	3.02%
(30, 0.6)	207,655	196,766	94.76%	212,704	188,936	88.83%	2.37%	4.14%
(30, 0.8)	167,445	158,680	94.77%	175,196	143,432	81.87%	4.42%	10.63%
Avg.			93.81%			84.43%	3.62%	7.44%

Table 2 Performance of the fluid approximations.

the same statistics for EXF. The eighth column gives the percent gap between the upper bounds provided by PRF and EXF, whereas the ninth column gives the percent gap between the total expected revenues of the approximate policies provided by the two fluid approximations.

Our results indicate that the approximate policy from PRF performs quite well. Over all of our test problems, the average gap between the total expected revenues of the approximate policy and the upper bounds provided by PRF is 6.19%. In other words, noting that we compare the performance of the approximate policy with an upper bound on the optimal total expected revenue, rather than the optimal total expected revenue itself, a conservative estimate of the average optimality gap of the approximate policy provided by PRF is 6.19%. In alignment with Theorem 4.1, the relative gaps between the upper bounds and the total expected revenues of the approximate policy from PRF tend to vanish as the number of stages increases. Considering groups of test problems with 5, 10, 15, 20, 25 and 30 stages, on average, the ratios between the total expected revenues of the approximate policy and the upper bounds provided by PRF are, respectively, 89.77, 92.89, 94.12, 95.47, 95.32 and 95.30. The performance of EXF, both in terms of the tightness of the upper bounds and the total expected revenues of the approximate policy that it provides, is noticeably inferior to PRF. Over all of our test problems, the average gap between the total expected revenues of




Figure 1 Performance of PRF on test problem with $K = 5$ for different values of the tuning parameter.

the approximate policy and the upper bounds provided by EXF is 15.57%. On average, the upper bounds on the optimal total expected revenue provided by PRF improve those provided by EXF by 3.62%. Similarly, the total expected revenues of the approximate policy provided by PRF improve those provided by EXF by 7.44%. There are test problems where the gap between the two upper bounds can reach 5.97% and the gap between the total expected revenues of the two approximate policies can reach 18.51%. Overall, EXF builds on the standard fluid approximation framework by using only the expected values of the demands, but there are significant benefits from using a fluid approximation that handles dependent demands with arbitrary distributions. In our computational experiments, we set the tuning parameter γ to one in our implementation of the approximate policy for PRF. While using a value of γ away from one is necessary to obtain our performance guarantee, our numerical experience indicates that using the value of one provides the strongest performance for our test problems. In Figure 1, we show the total expected revenue obtained by PRF as a function of γ for our problem instances with $K = 5$. The performance of our approximate policy is strongest when we set $\gamma = 1$. We observed a similar pattern in our other test problems.

Upper Bound from Problem (31): We give a problem instance to demonstrate that the upper bound provided by the optimal objective value of problem (31) is not asymptotically tight. In our problem instance, there is one stage, one resource and one product, so $K = 1$, $\mathcal{L} = \{1\}$ and $\mathcal{J} = \{1\}$. The capacity of the resource is α^2 for some $\alpha \in \mathbb{Z}_+$. The revenue of the product is $f_1 = 1$. The demand in the single stage has two possible values with $\mathbb{P}\{D^1 = \alpha\} = 1 - \frac{1}{2(\alpha-1)}$ and $\mathbb{P}\{D^1 = \alpha^3\} = \frac{1}{2(\alpha-1)}$. At each time period, there is a request for the product with probability one. For this problem instance, we have $K = 1$ and $c_{\min} = \alpha^2$. Furthermore, noting that there is one resource, we have $L = 1$. As a function of the parameter α in our problem instance, using OPT^α to denote the optimal total expected revenue, APX^α to denote the total expected revenue of the approximate policy and $\bar{Z}_{\text{LP}}^\alpha$ to denote the optimal objective value of problem (2), by Theorem 4.1, we have $\frac{\text{OPT}^\alpha}{\bar{Z}_{\text{LP}}^\alpha} \geq \frac{\text{APX}^\alpha}{\bar{Z}_{\text{LP}}^\alpha} \geq 1 - 4 \frac{\sqrt{\alpha^2 \log \alpha^2}}{\alpha^2} - \frac{1}{\alpha^2}$. By Theorem 3.1, we also have $1 \geq \frac{\text{OPT}^\alpha}{\bar{Z}_{\text{LP}}^\alpha}$. Thus, as α

becomes arbitrarily large, the relative gap between the optimal total expected revenue and the optimal objective value of problem (2) vanishes for our problem instance.

Using $\tilde{Z}_{\text{LP}}^\alpha$ to denote the optimal objective value of problem (31) as a function of the parameter α in our problem instance, we argue that $\frac{\text{OPT}^\alpha}{\tilde{Z}_{\text{LP}}^\alpha} = \frac{3}{\alpha+3}$. Therefore, as α becomes arbitrarily large, the relative gap between the optimal total expected revenue and the optimal objective value of problem (31) becomes arbitrarily large, so the optimal objective value of problem (31) is not asymptotically tight. Because there is a single product, it is optimal to accept the requests for the product as much as the capacity allows. If $D^1 = \alpha$, then the capacity of α^2 allows us to accept all product requests at α time periods and we obtain a total expected revenue of α . If $D^1 = \alpha^3$, then we can sell all of the capacity by using the product requests at α^3 time periods, so noting that the capacity of the resource is α^2 , we obtain a total expected revenue of α^2 . Thus, the optimal total expected revenue is $\text{OPT}^\alpha = \mathbb{P}\{D^1 = \alpha\} \alpha + \mathbb{P}\{D^1 = \alpha^3\} \alpha^2 = (1 - \frac{1}{2(\alpha-1)}) \alpha + \frac{1}{2(\alpha-1)} \alpha^2 = \frac{3\alpha}{2}$. On the other hand, for our problem instance, the linear program in (31) is given by

$$\tilde{Z}_{\text{LP}}^\alpha = \max_{w_1 \geq 0} \left\{ w_1 : w_1 \leq \alpha^2, \quad w_1 \leq \alpha + (\alpha^3 - \alpha) \frac{1}{2(\alpha - 1)} \right\}.$$

For $\alpha \geq 3$, we have $\alpha + (\alpha^3 - \alpha) \frac{1}{2(\alpha-1)} = \frac{\alpha^2+3\alpha}{2} \leq \alpha^2$, so $\tilde{Z}_{\text{LP}}^\alpha = \frac{\alpha^2+3\alpha}{2}$. Noting that $\text{OPT}^\alpha = \frac{3\alpha}{2}$, we obtain $\frac{\text{OPT}^\alpha}{\tilde{Z}_{\text{LP}}^\alpha} = \frac{3}{\alpha+3}$ for our problem instance, as long as $\alpha \geq 3$.