Online Appendix: Fluid Approximations for Revenue Management under High-Variance Demand

Yicheng Bai, Omar El Housni, Billy Jin, Paat Rusmevichientong, Huseyin Topaloglu, David P. Williamson

Appendix A: An Upper Bound on the Optimal Total Expected Revenue

In this section, we give a proof for Theorem 4.1. We will use an equivalent reformulation of the dynamic program in Section 2 that is more suitable for Lagrangian relaxation. In our equivalent reformulation, we use the decision variables $\boldsymbol{u} = (u_j : j \in \mathcal{J}) \in \{0,1\}^{|\mathcal{J}|}$, where $u_j = 1$ if and only if we make product j available at a generic time period. If the remaining capacities of the resources are given by the vector $\boldsymbol{x} \in \mathbb{Z}_+^{|\mathcal{L}|}$, then the set of feasible decisions is given by $\mathcal{U}(\boldsymbol{x}) = \{\boldsymbol{u} \in \{0,1\}^{|\mathcal{J}|} : \mathbf{1}_{(i \in A_j)}\, u_j \le x_i \ \forall\, i \in \mathcal{L},\ j \in \mathcal{J}\}$, which ensures that we can make product j available only when there is at least one unit of remaining capacity for all resources that are used by product j. In this case, the dynamic program in Section 2 is equivalent to

$$J_t(\boldsymbol{x}) = \max_{\boldsymbol{u} \in \mathcal{U}(\boldsymbol{x})} \left\{ \sum_{j \in \mathcal{J}} \lambda_{jt} \left\{ f_j\, u_j + \rho_t\, J_{t+1}\Big(\boldsymbol{x} - u_j \sum_{i \in A_j} \boldsymbol{e}_i\Big) \right\} \right\}, \tag{1}$$

with the boundary condition that $J_{T+1} = 0$. The value functions computed through the dynamic program in (1) are identical to those computed through the dynamic program in Section 2.

Considering (1), for each $i \in \mathcal{L}$ and $j \in \mathcal{J}$, we relax the constraint $\mathbf{1}_{(i \in A_j)}\, u_j \le x_i$ at time period t by associating the Lagrange multiplier θ_{ijt} with it to obtain the dynamic program

$$\begin{aligned} \hat{J}_t^{\boldsymbol{\theta}}(\boldsymbol{x}) &= \max_{\boldsymbol{u} \in \{0,1\}^{|\mathcal{J}|}} \left\{ \sum_{j \in \mathcal{J}} \lambda_{jt} \left\{ f_j\, u_j + \rho_t\, \hat{J}_{t+1}^{\boldsymbol{\theta}}\Big(\boldsymbol{x} - u_j \sum_{i \in A_j} \boldsymbol{e}_i\Big) \right\} + \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{J}} \lambda_{jt}\, \theta_{ijt} \Big[x_i - \mathbf{1}_{(i \in A_j)}\, u_j\Big] \right\} \\ &= \max_{\boldsymbol{u} \in \{0,1\}^{|\mathcal{J}|}} \left\{ \sum_{j \in \mathcal{J}} \lambda_{jt} \left\{ \Big[f_j - \sum_{i \in \mathcal{L}} \mathbf{1}_{(i \in A_j)}\, \theta_{ijt}\Big] u_j + \rho_t\, \hat{J}_{t+1}^{\boldsymbol{\theta}}\Big(\boldsymbol{x} - u_j \sum_{i \in A_j} \boldsymbol{e}_i\Big) \right\} \right\} + \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{J}} \lambda_{jt}\, \theta_{ijt}\, x_i, \end{aligned} \tag{2}$$

with the boundary condition that $\hat{J}_{T+1}^{\boldsymbol{\theta}} = 0$. In the first equality above, we scaled the Lagrange multiplier θ_{ijt} with λ_{jt}, which will simplify our notation. The second equality follows by arranging the terms. We refer to the dynamic program in (2) as the relaxed dynamic program. In the relaxed dynamic program, we make it explicit that the value functions depend on our choice of the Lagrange multipliers $\boldsymbol{\theta} = (\theta_{ijt} : i \in \mathcal{L},\ j \in \mathcal{J},\ t \in \mathcal{T}) \in \mathbb{R}_+^{|\mathcal{L}||\mathcal{J}||\mathcal{T}|}$. It is a standard result that the value functions of the relaxed dynamic program provide upper bounds on the value functions in (1); see, for example, Adelman and Mersereau (2008). Thus, for any choice of the Lagrange multipliers $\boldsymbol{\theta} \in \mathbb{R}_+^{|\mathcal{L}||\mathcal{J}||\mathcal{T}|}$, we have $\hat{J}_t^{\boldsymbol{\theta}}(\boldsymbol{x}) \ge J_t(\boldsymbol{x})$ for all $t \in \mathcal{T}$.

Noting that $J_1(\boldsymbol{c})$ is the optimal total expected revenue, by the discussion in the previous paragraph, for any $\boldsymbol{\theta} \in \mathbb{R}_+^{|\mathcal{L}||\mathcal{J}||\mathcal{T}|}$, $\hat{J}_1^{\boldsymbol{\theta}}(\boldsymbol{c})$ is an upper bound on the optimal total expected revenue. To show Theorem 4.1, we focus on a specific choice of the Lagrange multipliers, where the Lagrange multipliers at all time periods except for the last one are zero, whereas the Lagrange multipliers at the last time period depend on the resources but not on the products. For any $\boldsymbol{\eta} = (\eta_i : i \in \mathcal{L}) \in \mathbb{R}_+^{|\mathcal{L}|}$, define $\mathcal{F}(\boldsymbol{\eta}) = \{\boldsymbol{\theta} \in \mathbb{R}_+^{|\mathcal{L}||\mathcal{J}||\mathcal{T}|} : \theta_{ijT} = \eta_i \ \forall\, i \in \mathcal{L},\ j \in \mathcal{J},\ \theta_{ijt} = 0\ \forall\, i \in \mathcal{L},\ j \in \mathcal{J},\ t \in \mathcal{T} \setminus \{T\}\}$, so if $\boldsymbol{\theta} \in \mathcal{F}(\boldsymbol{\eta})$, then the Lagrange multipliers $(\theta_{ijT} : j \in \mathcal{J})$ take the common value of η_i, but the Lagrange multipliers $(\theta_{ijt} : i \in \mathcal{L}\ , j \in \mathcal{J},\ t \in \mathcal{T} \setminus \{T\})$ take the value of zero. Focusing on a specific choice of the Lagrange multipliers does not take full advantage of the flexibility provided by the possibility of using a different Lagrange multiplier for each resource, product and time period, but it will be enough to show that the optimal objective value of the Universal Fluid approximation is an upper bound on the optimal total expected revenue. In the next lemma, we give a closed form expression for the value functions $\{\hat{J}_t^{\boldsymbol{\theta}} : t \in \mathcal{T}\}$ when we have $\boldsymbol{\theta} \in \mathcal{F}(\boldsymbol{\eta})$.

Lemma A.1 *Letting $R_t = \rho_t\, \rho_{t+1} \ldots \rho_{T-1}$ with $R_T = 1$, for any $\boldsymbol{\eta} \in \mathbb{R}_+^{|\mathcal{L}|}$, if the Lagrange multipliers satisfy $\boldsymbol{\theta} \in \mathcal{F}(\boldsymbol{\eta})$, then we have*

$$\hat{J}_t^{\boldsymbol{\theta}}(\boldsymbol{x}) = \sum_{k=t}^{T} \sum_{j \in \mathcal{J}} \lambda_{jk} \Big[\frac{R_t}{R_k} f_j - R_t \sum_{i \in A_j} \eta_i \Big]^+ + R_t \sum_{i \in \mathcal{L}} \eta_i\, x_i.$$

Proof: We show the result by using induction over the time periods. At time period T, noting that $\hat{J}_{T+1}^{\boldsymbol{\theta}} = 0$ and $\theta_{ijT} = \eta_i$, as well as using the fact that $\sum_{j \in \mathcal{J}} \lambda_{jT} = 1$, by (2), we have

$$\hat{J}_T^{\boldsymbol{\theta}}(\boldsymbol{x}) = \max_{\boldsymbol{u} \in \{0,1\}^{|\mathcal{J}|}} \Bigg\{ \sum_{j \in \mathcal{J}} \lambda_{jT} \Big[f_j - \sum_{i \in \mathcal{L}} \mathbf{1}_{(i \in A_j)}\, \eta_i \Big] u_j \Bigg\} + \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{J}} \lambda_{jT}\, \eta_i\, x_i = \sum_{j \in \mathcal{J}} \lambda_{jT} \Big[f_j - \sum_{i \in A_j} \eta_i \Big]^+ + \sum_{i \in \mathcal{L}} \eta_i\, x_i.$$

Thus, the result holds at time period T. Assuming that the result holds at time period $t+1$, we show that the result holds at time period t as well.

Letting $K_t = \sum_{k=t}^{T} \sum_{j \in \mathcal{J}} \lambda_{jk} \big[\frac{R_t}{R_k} f_j - R_t \sum_{i \in A_j} \eta_i \big]^+$, by the induction assumption, we have $\hat{J}_{t+1}^{\boldsymbol{\theta}}(\boldsymbol{x}) = K_{t+1} + R_{t+1} \sum_{i \in \mathcal{L}} \eta_i\, x_i$. Noting that $\theta_{ijt} = 0$ for $t \in \mathcal{T} \setminus \{T\}$, by (2), we have

$$\begin{aligned}
\hat{J}_t^{\boldsymbol{\theta}}(\boldsymbol{x}) &\overset{(a)}{=} \max_{\boldsymbol{u} \in \{0,1\}^{|\mathcal{J}|}} \Bigg\{ \sum_{j \in \mathcal{J}} \lambda_{jt} \Big\{ f_j\, u_j + \rho_t\, K_{t+1} + \rho_t\, R_{t+1} \sum_{i \in \mathcal{L}} \eta_i\, (x_i - \mathbf{1}_{(i \in A_j)}\, u_j) \Big\} \Bigg\} \\
&\overset{(b)}{=} \max_{\boldsymbol{u} \in \{0,1\}^{|\mathcal{J}|}} \Bigg\{ \sum_{j \in \mathcal{J}} \lambda_{jt} \Big[f_j - \rho_t\, R_{t+1} \sum_{i \in A_j} \eta_i \Big] u_j \Bigg\} + \rho_t\, K_{t+1} + \rho_t\, R_{t+1} \sum_{i \in \mathcal{L}} \eta_i\, x_i \\
&= \sum_{j \in \mathcal{J}} \lambda_{jt} \Big[f_j - \rho_t\, R_{t+1} \sum_{i \in A_j} \eta_i \Big]^+ + \rho_t\, K_{t+1} + \rho_t\, R_{t+1} \sum_{i \in \mathcal{L}} \eta_i\, x_i \\
&\overset{(c)}{=} K_t + R_t \sum_{i \in \mathcal{L}} \eta_i\, x_i,
\end{aligned}$$

where (a) holds because we have $\hat{J}_{t+1}^{\boldsymbol{\theta}}(\boldsymbol{x} - u_j \sum_{i \in A_j} \boldsymbol{e}_i) = K_{t+1} + R_{t+1} \sum_{i \in \mathcal{L}} \eta_i\, (x_i - \mathbf{1}_{(i \in A_j)}\, u_j)$ by the induction assumption, (b) follows by arranging the terms and using the fact that $\sum_{j \in \mathcal{J}} \lambda_{jt} = 1$

and (c) holds because we have $R_t = \rho_t \, R_{t+1}$ and $K_t = \sum_{j \in \mathcal{J}} \lambda_{jt} \, [f_j - R_t \sum_{i \in A_j} \theta_i]^+ + \rho_t \, K_{t+1}$ by the definitions of R_t and K_t. By the chain of equalities above, the result holds at time period t. ■

We have $R_t = \rho_t \, \rho_{t+1} \ldots \rho_{T-1} = \frac{\mathbb{P}\{D \geq t+1\}}{\mathbb{P}\{D \geq t\}} \frac{\mathbb{P}\{D \geq t+2\}}{\mathbb{P}\{D \geq t+1\}} \cdots \frac{\mathbb{P}\{D \geq T\}}{\mathbb{P}\{D \geq T-1\}} = \frac{\mathbb{P}\{D \geq T\}}{\mathbb{P}\{D \geq t\}}$. Since $\mathbb{P}\{D \geq 1\} = 1$, the last equality also yields $R_1 = \mathbb{P}\{D \geq T\}$. By Lemma A.1, for any $\boldsymbol{\eta} \in \mathbb{R}_+^{|\mathcal{L}|}$ and $\boldsymbol{\theta} \in \mathcal{F}(\boldsymbol{\eta})$, we get

$$\begin{aligned} \hat{J}_1^{\boldsymbol{\theta}}(\boldsymbol{c}) \; &= \; \sum_{t=1}^{T} \sum_{j \in \mathcal{J}} \lambda_{jt} \left[\frac{R_1}{R_t} f_j - R_1 \sum_{i \in A_j} \eta_i \right]^+ + R_1 \sum_{i \in \mathcal{L}} \eta_i \, c_i \\ &= \; \sum_{t=1}^{T} \sum_{j \in \mathcal{J}} \lambda_{jt} \left[\mathbb{P}\{D \geq t\} \, f_j - \mathbb{P}\{D \geq T\} \sum_{i \in \mathcal{L}} \mathbf{1}_{(i \in A_j)} \, \eta_i \right]^+ + \; \mathbb{P}\{D \geq T\} \sum_{i \in \mathcal{L}} \eta_i \, c_i. \end{aligned} \quad (3)$$

We have $\hat{J}_1^{\boldsymbol{\theta}}(\boldsymbol{c}) \geq J_1(\boldsymbol{c})$ for all $\boldsymbol{\theta} \in \mathbb{R}_+^{|\mathcal{L}||\mathcal{J}||\mathcal{T}|}$, but we have $\boldsymbol{\theta} \in \mathbb{R}_+^{|\mathcal{L}||\mathcal{J}||\mathcal{T}|}$ for any $\boldsymbol{\eta} \in \mathbb{R}_+^{|\mathcal{L}|}$ and $\boldsymbol{\theta} \in \mathcal{F}(\boldsymbol{\eta})$. Thus, the expression on the right side of (3) is an upper bound on $J_1(\boldsymbol{c})$ for any $\boldsymbol{\eta} \in \mathbb{R}_+^{|\mathcal{L}|}$.

By the discussion in the previous paragraph, if we minimize the expression on the right side of (3) over all $\boldsymbol{\eta} \in \mathbb{R}_+^{|\mathcal{L}|}$, then we get an upper bound on $J_1(\boldsymbol{c})$. Below is the proof of Theorem 4.1.

Proof of Theorem 4.1:

Using the decision variables $\boldsymbol{\eta} = (\eta_i : i \in \mathcal{L}) \in \mathbb{R}_+^{|\mathcal{L}|}$ and $\boldsymbol{z} = (z_{jt} : j \in \mathcal{J}, \; t \in \mathcal{T}) \in \mathbb{R}_+^{|\mathcal{J}||\mathcal{T}|}$, we can minimize the expression on the right side of (3) over all $\boldsymbol{\eta} \in \mathbb{R}_+^{|\mathcal{L}|}$ by solving the linear program

$$\min_{(\boldsymbol{\eta}, \boldsymbol{z}) \in \mathbb{R}_+^{|\mathcal{L}| + |\mathcal{J}||\mathcal{T}|}} \left\{ \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \lambda_{jt} \, z_{jt} + \mathbb{P}\{D \geq T\} \sum_{i \in \mathcal{L}} c_i \, \eta_i \; : \; z_{jt} \; \geq \; \mathbb{P}\{D \geq t\} \, f_j \; - \; \mathbb{P}\{D \geq T\} \sum_{i \in \mathcal{L}} \mathbf{1}_{(i \in A_j)} \, \eta_i \quad \forall \, j \in \mathcal{J}, \; t \in \mathcal{T} \right\}. \quad (4)$$

We can assume that $\mathbb{P}\{D \geq T\} > 0$, because if $\mathbb{P}\{D \geq T\} = 0$, then we can choose the upper bound of the support of D as the largest value of $\tau \in \{1, \ldots, T-1\}$ such that $\mathbb{P}\{D \geq \tau\} > 0$.

Associating the dual variables $\boldsymbol{y} = (y_{jt} : j \in \mathcal{J}, \; t \in \mathcal{T}) \in \mathbb{R}_+^{|\mathcal{J}||\mathcal{T}|}$ with the constraints in the problem above, the dual of problem (4) is given by

$$\max_{\boldsymbol{y} \in \mathbb{R}_+^{|\mathcal{J}||\mathcal{T}|}} \left\{ \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} f_j \, \mathbb{P}\{D \geq t\} \, y_{jt} \;\; : \;\; \mathbb{P}\{D \geq T\} \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \mathbf{1}_{(i \in A_j)} \, y_{jt} \; \leq \; \mathbb{P}\{D \geq T\} \, c_i \quad \forall \, i \in \mathcal{L} \quad y_{jt} \; \leq \; \lambda_{jt} \quad \forall \, j \in \mathcal{J}, \; t \in \mathcal{T} \right\}. \quad (5)$$

Because $\mathbb{P}\{D \geq T\} > 0$, problem (5) is equivalent to the Universal Fluid. The desired result follows since (4) and (5) have the same optimal objective value, which is an upper bound on $J_1(\boldsymbol{c})$. ■

In our derivation of the Universal Fluid approximation, we used the relaxed dynamic program with a specific choice of the Lagrange multipliers, where the Lagrange multipliers at all time periods except for the last one are zero, whereas the Lagrange multipliers at the last time period depend on the resources but not on the products. We can potentially derive other fluid approximations by using more general Lagrange multipliers that exploit the possibility that we can use a different Lagrange multiplier for each resource, product and time period. The upper bounds on the optimal total expected revenue provided by such fluid approximations would be at least as tight as those provided by the Universal Fluid approximation, but such fluid approximations do not have a form that is as simple as the Universal Fluid approximation. Furthermore, the Universal Fluid approximation turns out to be enough to obtain asymptotically tight upper bounds.

Appendix B: Asymptotic Tightness of the Fluid Approximation

We give a proof for Theorem 5.1. Let $\boldsymbol{y}^* = (y^*_{jt} : j \in \mathcal{J},\ t \in \mathcal{T})$ be an optimal solution to the Universal Fluid approximation. Consider the following approximate policy for some $\theta \in (0,1)$. At time period t, we make product j available for purchase with probability $\theta \frac{y^*_{jt}}{\lambda_{jt}}$. If the customer arriving at time period t wants to purchase product j and there is capacity available to serve a request for product j, then we sell a unit of product j and consume the capacities of the resources used by the product. Define three Bernoulli random variables. The first one, denoted by A_{jt}, takes value one if the approximate policy makes product j available at time period t. We have $\mathbb{P}\{A_{jt} = 1\} = \theta \frac{y^*_{jt}}{\lambda_{jt}}$. The second one, denoted by Λ_{jt}, takes value one if the customer arriving at time period t is interested in purchasing product j. We have $\mathbb{P}\{\Lambda_{jt} = 1\} = \lambda_{jt}$. The third one, denoted by G_{jt}, takes value one if we have capacity to serve a request for product j at time period t under the approximate policy. In this case, the total revenue obtained by the approximate policy is given by the random variable $\sum_{t=1}^{T} \sum_{j \in \mathcal{J}} f_j \, \mathbf{1}_{(D \geq t,\, A_{jt}=1,\, \Lambda_{jt}=1,\, G_{jt}=1)}$, where we use the fact that the approximate policy makes a sale for product j at time period t if and only if the selling horizon reaches beyond this time period, the approximate policy makes product j available, the arriving customer is interested in purchasing product j and we have capacity to serve a request for product j. Note that G_{jt} depends on the decisions of the approximate policy at time periods $1, \ldots, t-1$ and D is independent of the decisions of the approximate policy and the products of interest to the arriving customers. Thus, letting $J_{\mathsf{App}}(\boldsymbol{c})$ be the total expected revenue of the approximate policy, we get

$$J_{\mathsf{App}}(\boldsymbol{c}) = \sum_{t=1}^{T} \sum_{j \in \mathcal{J}} f_j \, \mathbb{P}\{D \geq t\} \, \mathbb{P}\{A_{jt} = 1\} \, \mathbb{P}\{\Lambda_{jt} = 1\} \, \mathbb{P}\{G_{jt} = 1\}$$

$$= \sum_{t=1}^{T} \sum_{j \in \mathcal{J}} f_j \, \mathbb{P}\{D \geq t\} \, \theta \, \frac{y^*_{jt}}{\lambda_{jt}} \, \lambda_{jt} \, \mathbb{P}\{G_{jt} = 1\}. \quad (6)$$

We lower bound the probability $\mathbb{P}\{G_{jt} = 1\}$. At time period t, the approximate policy makes product j available with probability $\theta \frac{y^*_{jt}}{\lambda_{jt}}$, whereas we have a request for product j with probability λ_{jt}. Thus, under the approximate policy, there is a unit of demand for capacity of resource i at time period t with probability $\sum_{j\in\mathcal{J}} \mathbf{1}_{(i\in A_j)}\,\theta\,\frac{y^*_{jt}}{\lambda_{jt}}\,\lambda_{jt} = \theta \sum_{j\in\mathcal{J}} \mathbf{1}_{(i\in A_j)}\,y^*_{jt}$. However, having demand for capacity of resource i at time period t does not mean that the approximate policy depletes the capacity of the resource at time period t. In particular, considering some product j that uses the capacity of resource i, even if the approximate policy makes product j available at time period t and the customer arriving at time period t is interested in product j, we may not have capacity on some other resource used by product j, in which case, we would not be serving the demand for the product. Thus, letting $\{N_{it} : t\in\mathcal{T}\}$ be a collection of independent Bernoulli random variables, where N_{it} takes value one with probability $\theta\sum_{j\in\mathcal{J}} \mathbf{1}_{(i\in A_j)}\,y^*_{jt}$, under the approximate policy, the total capacity consumption of resource i over time periods $1,\ldots,t-1$ is upper bounded by $\sum_{k=1}^T N_{ik}$. Thus, having $\sum_{k=1}^T N_{ik} < c_i$ for all $i\in A_j$ implies that $G_{jt}=1$, so $\mathbb{P}\{\sum_{k=1}^T N_{ik} < c_i \ \ \forall\, i\in A_j\} \le \mathbb{P}\{G_{jt}=1\}$. We need the concentration bound in the next lemma.

Lemma B.1 *Letting $\{N_{it} : t\in\mathcal{T}\}$ be a collection of independent Bernoulli random variables, where N_{it} takes value one with probability $\theta\sum_{j\in\mathcal{J}} \mathbf{1}_{(i\in A_j)}\,y^*_{jt}$, we have*

$$\mathbb{P}\left\{\sum_{t=1}^T N_{it} \ge c_i\right\} \ \le\ \exp\left(-\frac{\frac{3}{2}(1-\theta)^2 c_{\min}}{2\theta+1}\right).$$

Proof: Letting $\rho_{it} = \theta\sum_{j\in\mathcal{J}} \mathbf{1}_{(i\in A_j)}\,y^*_{jt}$ for notational brevity, so that we have $\mathbb{E}\{N_{it}\} = \rho_{it}$ and $\mathrm{Var}(N_{it}) = \rho_{it}\,(1-\rho_{it})$, we upper bound the expectation and variance of $\sum_{t=1}^T N_{it}$ as

$$\mathrm{Var}\left(\sum_{t=1}^T N_{it}\right) = \sum_{t=1}^T \rho_{it}\,(1-\rho_{it}) \le \sum_{t=1}^T \rho_{it} = \mathbb{E}\left\{\sum_{t=1}^T N_{it}\right\} = \theta \sum_{t=1}^T\sum_{j\in\mathcal{J}} \mathbf{1}_{(i\in A_j)}\,y^*_{jt} \le \theta\, c_i,$$

where the last inequality holds because $\boldsymbol{y}^* = (y^*_{jt} : j\in\mathcal{J},\ t\in\mathcal{T})$ is an optimal solution to the Universal Fluid approximation, so it satisfies the first constraint in the fluid approximation.

Noting that $\mathbb{E}\{\sum_{t=1}^T N_{it}\} \le \theta\, c_i$ by the chain of inequalities above, using the one-sided Bernstein inequality, we obtain the chain of inequalities

$$\begin{aligned}\mathbb{P}\left\{\sum_{t=1}^T N_{it} \ge c_i\right\} &\overset{(a)}{\le} \mathbb{P}\left\{\sum_{t=1}^T \Big[N_{it} - \mathbb{E}\{N_{it}\}\Big] \ge (1-\theta)\,c_i\right\} \overset{(b)}{\le} \exp\left(-\frac{\frac{1}{2}(1-\theta)^2\,c_i^2}{\sum_{t=1}^T \mathrm{Var}(N_{it}) + \frac{1}{3}(1-\theta)c_i}\right)\\ &\overset{(c)}{\le} \exp\left(-\frac{\frac{1}{2}(1-\theta)^2\,c_i^2}{\theta\,c_i + \frac{1}{3}(1-\theta)c_i}\right) = \exp\left(-\frac{\frac{3}{2}(1-\theta)^2\,c_i}{2\theta+1}\right) \overset{(d)}{\le} \exp\left(-\frac{\frac{3}{2}(1-\theta)^2\,c_{\min}}{2\theta+1}\right),\end{aligned}$$

where (a) holds because $\mathbb{E}\{\sum_{t=1}^T N_{it}\} \le \theta\, c_i$, (b) is the one-sided Bernstein inequality, (c) uses the fact that $\sum_{t=1}^T \mathrm{Var}(N_{it}) \le \theta\, c_i$ and (d) uses the fact that $c_{\min} \le c_i$. ■

We can use the bound in the lemma above along with the union bound to come up with a lower bound on the probability $\mathbb{P}\{\sum_{t=1}^T N_{it} < c_i \ \ \forall\, i\in A_j\}$. By the discussion right before the lemma, a lower bound on the last probability is also a lower bound on the probability $\mathbb{P}\{G_{jt}=1\}$. Putting these observations together will yield a proof for Theorem 5.1.

Proof of Theorem 5.1:

Noting the discussion just before Lemma B.1, $\mathbb{P}\{\sum_{k=1}^T N_{ik} < c_i \;\; \forall i \in A_j\} \le \mathbb{P}\{G_{jt}=1\}$. We lower bound the probability $\mathbb{P}\{G_{jt}=1\}$ as

$$\mathbb{P}\{G_{jt}=1\} \ge \mathbb{P}\left\{\sum_{t=1}^T N_{it} < c_i \;\; \forall i \in A_j\right\} = 1 - \mathbb{P}\left\{\exists i \in A_j \text{ such that } \sum_{t=1}^T N_{it} \ge c_i \right\}$$
$$\overset{(a)}{\ge} 1 - \sum_{i\in A_j} \mathbb{P}\left\{\sum_{t=1}^T N_{it} \ge c_i \right\} \overset{(b)}{\ge} 1 - L\exp\left(-\frac{\frac{3}{2}(1-\theta)^2 c_{\min}}{2\theta+1}\right) \overset{(c)}{\ge} 1 - L\exp\left(-\frac{(1-\theta)^2 c_{\min}}{2}\right),$$

where (a) is the union bound, (b) follows from Lemma B.1, as well as the fact that $|A_j| \le L$ and (c) uses the fact that $\theta \in (0,1)$, in which case, we have $2\theta+1<3$.

If we use $\theta = 1 - \sqrt{\frac{2\log c_{\min}}{c_{\min}}}$ in our approximate policy, then the right of the chain of inequalities above reads $1-\frac{L}{c_{\min}}$, so $\mathbb{P}\{G_{jt}=1\} \ge 1-\frac{L}{c_{\min}}$ with this choice of θ. Thus, by (6), we get

$$J_{\mathsf{App}}(\boldsymbol{c}) \;\ge\; \left(1-\sqrt{\frac{2\log c_{\min}}{c_{\min}}}\right)\sum_{t=1}^T\sum_{j\in\mathcal{J}} f_j\,\mathbb{P}\{D\ge t\}\,y^*_{jt}\left(1-\frac{L}{c_{\min}}\right)$$
$$\overset{(d)}{=} \left(1-\sqrt{\frac{2\log c_{\min}}{c_{\min}}}\right)\left(1-\frac{L}{c_{\min}}\right)Z^*_{\mathsf{LP}} \;\ge\; \left(1-\sqrt{\frac{2\log c_{\min}}{c_{\min}}}-\frac{L}{c_{\min}}\right)Z^*_{\mathsf{LP}},$$

where (d) holds because the solution $\boldsymbol{y}^* = (y^*_{jt} : j\in\mathcal{J},\; t\in\mathcal{T})$ is optimal to the Universal Fluid. The desired result follows because the optimal total expected revenue satisfies $J_1(\boldsymbol{c}) \ge J_{\mathsf{App}}(\boldsymbol{c})$. ■

Appendix C: Experimental Setup for the Test Problems

We give the details of our approach for generating our test problems. Letting Γ be a log-normal random variable with mean μ and standard deviation μv and k be the smallest integer such that $\mathbb{P}\{\Gamma \le k\} \ge 0.99$, we set the maximum length of the selling horizon as $T=k$. For each $t=1,\ldots,T$, letting $\gamma_t = \mathbb{P}\{t-1\le\Gamma\le t\}$, the probability mass function of D evaluated at t is proportional to γ_t. In particular, for each $t=1,\ldots,T$, we set $\mathbb{P}\{D=t\} = \gamma_t/\sum_{s=1}^T \gamma_s$. We place the hub at the center of a 100×100 square and generate the locations of the spokes uniformly over the same square. The fare associated with a low-fare itinerary is the sum of the Euclidean distances traversed by the flights in the itinerary. The fare associated with a high-fare itinerary is κ times the fare of the corresponding low-fare itinerary. We set $\kappa=4$.

To come up with the arrival probabilities for the customers interested in different itineraries, for each origin-destination pair (f,g), we generate ξ_{fg} from the uniform distribution over $[0,1]$.

One of the locations in the origin-destination pair can be the hub. Letting N be the set of all locations, we normalize these samples by setting $\zeta_{fg} = \xi_{fg} / \sum_{(p,q)\in N^2, p\neq q} \xi_{pq}$ so that they add up to one. The probability that the customer arriving at any time period is interested in an itinerary that connects the origin-destination pair (f,g) is ζ_{fg}. The probability that a customer is interested in a low-fare itinerary decreases over time, whereas we have the reverse trend for the probability that a customer is interested in a high-fare itinerary. In this way, we generate test problems where the requests for high-fare itineraries tend to arrive later and we need to protect capacity for the high-fare itinerary requests that tend to arrive later. To generate test problems with this feature, for each origin-destination pair (f,g), we generate a time threshold τ_{fg} uniformly over $\{1,\ldots,T\}$. The probability of having a request for a low-fare itinerary linearly decreases over time, whereas the probability of having a request for a high-fare itinerary is zero until time period τ_{fg}, but it increases linearly after time period τ_{fg}. In particular, we define the functions $G, H_{fg} : \mathcal{T} \to \mathbb{R}_+$ as $G(t) = 1 - \frac{t-1}{T-1}$ and $H_{fg}(t) = \left[\frac{t-\tau_{fg}}{T-\tau_{fg}}\right]^+$. In this case, if itinerary j is the low-fare itinerary connecting origin-destination pair (f,g), then $\lambda_{jt} = \zeta_{fg} \frac{G(t)}{G(t)+H_{fg}(t)}$ and if itinerary j is the high-fare itinerary connecting origin-destination pair (f,g), then $\lambda_{jt} = \zeta_{fg} \frac{H_{fg}(t)}{G(t)+H_{fg}(t)}$. Once we generate the customer arrival probabilities, we set the capacities of the flight legs so that the total expected demand for the capacity on the flight leg exceeds the capacity of the flight leg by a factor of 1.6. In other words, noting that the total expected demand for the capacity on flight leg i is $\sum_{t\in\mathcal{T}} \sum_{j\in\mathcal{J}} \mathbf{1}_{(i\in A_j)} \lambda_{jt}$, the capacity of flight leg i is $c_i = \lceil \sum_{t\in\mathcal{T}} \sum_{j\in\mathcal{J}} \mathbf{1}_{(i\in A_j)} \lambda_{jt}/1.6 \rceil$.

We can choose the coefficient of variation of a log-normal random variable as large as we would like, which was the motivation for using this distribution for D in our experimental setup.

Appendix D: Finite Upper Bound on the Number of Customer Arrivals

We start by considering the case where there exists some $\overline{\lambda} > 0$ such that $\lambda_{jt} \geq \overline{\lambda}$ for all $j \in \mathcal{J}$ and $t \in \mathcal{T}$. Thus, the probability that an arriving customer is interested in a particular product is uniformly lower bounded by $\overline{\lambda}$. We discuss relaxations of this setup at the end of this section. Making the dependence of the Universal Fluid on the set of possible values for the length of the selling horizon explicit, we write the optimal objective value of this problem as $Z^*_{\mathsf{LP}}(\mathcal{T})$. Define the time threshold $\tau = \lceil \max_{i\in\mathcal{L}} c_i / \overline{\lambda} \rceil$. In the next proposition, we show that if $T > \tau$, then we can drop the last time period T in the Universal Fluid approximation. Therefore, we can always use a finite upper bound of τ on the possible realizations of the number of customer arrivals.

Proposition D.1 *If $T > \tau$, then we have $Z^*_{\mathsf{LP}}(\mathcal{T}) = Z^*_{\mathsf{LP}}(\mathcal{T} \setminus \{T\})$.*

Proof: Let $\boldsymbol{y}^* = (y^*_{jt} : j \in \mathcal{J},\ t \in \mathcal{T})$ be an optimal solution to the Universal Fluid. If $y^*_{jT} = 0$ for all $j \in \mathcal{J}$, then the result follows. Otherwise, there exists some product k such that $y^*_{kT} > 0$. We will construct another optimal solution $\widehat{\boldsymbol{y}} = (\hat{y}_{jt} : j \in \mathcal{J},\ t \in \mathcal{T})$ to the Universal Fluid approximation such that $\hat{y}_{kT} = 0$ and $\hat{y}_{jT} = y^*_{jtT}$ for all $j \in \mathcal{J} \setminus \{k\}$, in which case, repeatedly applying the same construction for each $k \in \mathcal{J}$ such that $y^*_{kT} > 0$, the desired result follows. Let $k \in \mathcal{J}$ be such that $y^*_{kT} > 0$. Choose some resource i that is used by product k, so $\mathbf{1}_{(i \in A_k)} = 1$. Using the fact that $T > \tau$ we have $y^*_{kT} + \sum_{t=1}^{\tau} y^*_{kt} \le \sum_{t \in \mathcal{T}} \mathbf{1}_{(i \in A_k)}\, y^*_{kt} \le \sum_{j \in \mathcal{J}} \sum_{t \in \mathcal{T}} \mathbf{1}_{(i \in A_j)}\, y^*_{jt} \le c_i$, where the last inequality holds because $\boldsymbol{y}^*$ satisfies the first constraint in the Universal Fluid approximation. By the definition of τ, we have $\tau \ge c_i / \overline{\lambda}$, so the last chain of inequalities yields $y^*_{kT} + \sum_{t=1}^{\tau} y^*_{kT} \le \tau\, \overline{\lambda} \le \sum_{t=1}^{\tau} \lambda_{kt}$, where we use the fact that $\lambda_{kt} \ge \overline{\lambda}$ for all $t \in \mathcal{T}$. Thus, we have $y^*_{kT} \le \sum_{t=1}^{\tau} (\lambda_{kt} - y^*_{kt})$. Noting that $\lambda_{kt} - y^*_{kt} \ge 0$ for all $t = 1, \ldots, \tau$ by the second constraint in the Universal Fluid approximation, having $y^*_{kT} \le \sum_{t=1}^{\tau} (\lambda_{kt} - y^*_{kt})$ implies that there exists a collection of non-negative numbers $\delta_1, \ldots, \delta_\tau$ such that we have $\sum_{t=1}^{\tau} \delta_t = y^*_{kT}$ and $\delta_t \le \lambda_{kt} - y^*_{kt}$ for all $t = 1, \ldots, \tau$. In this case, we define the solution $\widehat{\boldsymbol{y}} = (\hat{y}_{jt} : j \in \mathcal{J},\ t \in \mathcal{T})$ as $\hat{y}_{jt} = y^*_{jt}$ for all $j \in \mathcal{J} \setminus \{k\}$, $t \in \mathcal{T}$ and

$$\hat{y}_{kt} = \begin{cases} y^*_{kt} + \delta_t & \text{if } t = 1, \ldots, \tau \\ y^*_{kt} & \text{if } t = \tau + 1, \ldots, T - 1 \\ 0 & \text{if } t = T. \end{cases}$$

Because $\sum_{t=1}^{\tau} \delta_t = y^*_{kT}$, we have $\sum_{t \in \mathcal{T}} \hat{y}_{jt} = \sum_{t \in \mathcal{T}} y^*_{jt}$ for all $j \in \mathcal{J}$, so noting that $\boldsymbol{y}^*$ satisfies the first constraint in the Universal Fluid approximation, $\widehat{\boldsymbol{y}}$ satisfies this constraint too. Because $\delta_t \le \lambda_{kt} - y^*_{kt}$, the solution $\widehat{\boldsymbol{y}}$ satisfies the second constraint in the Universal Fluid approximation as well.

Thus, the solution $\widehat{\boldsymbol{y}}$ is feasible to the Universal Fluid approximation. The difference between the objective function values provided by $\widehat{\boldsymbol{y}}$ and $\boldsymbol{y}^*$ is $f_k \sum_{t=1}^{\tau} \mathbb{P}\{D \ge t\}\, \delta_t - f_k\, \mathbb{P}\{D \ge T\}\, y^*_{kT} = f_k \sum_{t=1}^{\tau} [\mathbb{P}\{D \ge t\} - \mathbb{P}\{D \ge T\}]\, \delta_t \ge 0$, where we use $\mathbb{P}\{D \ge t\} \ge \mathbb{P}\{D \ge T\}$ for all $t = 1, \ldots, \tau$. ■

By the proposition above, we can drop all time periods in $\mathcal{T} \setminus \{1, \ldots, \tau\}$ from consideration in the Universal Fluid approximation. We can extend the proposition above to the case where there exists some $\overline{\lambda} > 0$ such that $\lambda_{jt} \ge \mathbf{1}_{(\lambda_{jt} > 0)}\, \overline{\lambda}$ for all $j \in \mathcal{J}$ and $t \in \mathcal{T}$, so that the nonzero values for the probability that an arriving customer is interested in a particular product is uniformly lower bounded by $\overline{\lambda}$. In this case, we define as τ before. Furthermore, for each $j \in \mathcal{J}$, we choose $K_j = 1, \ldots, T + 1$ such that we have either $\sum_{t=1}^{K_j} \mathbf{1}_{(\lambda_{jt} > 0)} \ge \tau$ or $\sum_{K_j}^{T} \mathbf{1}_{(\lambda_{jt} > 0)} = 0$. Note that we can always choose $K_j = T + 1$, so there is always a value for K_j that satisfies one of the two conditions. In this case, we can show that if $T > \max_{j \in \mathcal{J}} K_j$, then we have $Z^*_{\mathsf{LP}}(\mathcal{T}) = Z^*_{\mathsf{LP}}(\mathcal{T} \setminus \{T\})$. In particular, if $\sum_{K_j}^{T} \mathbf{1}_{(\lambda_{jt} > 0)} = 0$, then there are no requests for product j after time period K_j. Thus, we can indeed set the decision variable y_{jT} to zero in the Universal Fluid. On the other hand, if $\sum_{t=1}^{K_j} \mathbf{1}_{(\lambda_{jt} > 0)} \ge \tau$, then there are τ time periods at which there is demand for product j with a probability of at least $\overline{\lambda}$. In this case, we can use the same argument in the proof of Proposition D.1 to conclude that we can set the decision variable y_{jT} to zero in the Universal Fluid.

References

Adelman, D., A. J. Mersereau. 2008. Relaxations of weakly coupled stochastic dynamic programs. *Operations Research* **56**(3) 712–727.