Appendix A: Missing Proofs - Section 3

A.1. Theorem 1 - Proof of negative result

Here, we establish the hardness result provided in Theorem 1. Namely, we show that no policy can achieve a competitive ratio better than ≈ 0.853. To do, we employ Yao's minimax principle, which in our context, stipulates that

$$\max_{\pi\in\Pi_D} \mathbb{E}_{S\sim X}\left[\frac{R(\mathcal{A}_\pi(S))}{\mathrm{OPT}(S)}\right] \geq \max_{\pi\in\Pi} \mathrm{CR}_\pi,$$

where Π_D is the collection of all admissible deterministic policies, Π is the collection of the policies that randomize over any deterministic policy $\pi_D \in \Pi_D$, and X is a distribution over deposit sequences. We therefore proceed to specify a distribution X for which

$$0.853 \geq \max_{\pi\in\Pi_D} \mathbb{E}_{S\sim X}\left[\frac{R(\mathcal{A}_\pi(S))}{\mathrm{OPT}(S)}\right].$$

Namely, consider the following two-sequence distribution, where $\Pr[X = S_1] = \alpha$ and $\Pr[X = S_2] = 1-\alpha$, and two deposit sequences are defined as follow:

- The sequence S_1 is comprised of m identical copies of the tuple $(p_t, d_t) = (0.413, 0.707)$, noting that if S_1 is realized, then we clearly have that $\mathrm{OPT}(S_1) = m\cdot(0.413 + (1-0.413)\cdot 0.707) = 0.828m$. These parameter values were reverse engineered to give the strongest impossibility result.
- The other sequence S_2 mirrors S_1 for the first m shippers, who are followed by m subsequent shippers with a deposit tuple of $(1,1)$. If S_2 is realized, it is clearly optimal for the clairvoyant to reject the first m shippers, while accepting the latter m shipper, leading to a profit of $\mathrm{OPT}(S_2) = m$.

Given the two-sequence distribution specified above, it is straightforward to see that any deterministic policy $\pi \in \Pi_D$ can be defined by specifying solely the number of shippers $k \in [m]_0 = \{0, 1, 2, \ldots, m\}$ accepted among the first m arrivals. From here, if sequence S_2 is realized, we have crafted the parameters of the two sequences so that it is optimal to accept $m-k$ shippers from the second batch. In turn, by choosing the worst-case value of α, we get that

$$\begin{aligned}
\min_{\alpha\in[0,1]}\max_{\pi\in\Pi_D} \mathbb{E}_{S\sim X}\left[\frac{R(\mathcal{A}_\pi(S))}{\mathrm{OPT}(S)}\right] &= \min_{\alpha\in[0,1]}\max_{k\in[m]_0} \alpha\frac{0.828k}{\mathrm{OPT}(S_1)} + (1-\alpha)\cdot\frac{0.828k + (m-k)}{\mathrm{OPT}(S_2)} \\
&= \min_{\alpha\in[0,1]}\max_{k\in[m]_0} \alpha\frac{k}{m} + (1-\alpha)\cdot\frac{0.828k + (m-k)}{m} \\
&\leq \min_{\alpha\in[0,1]}\max_{\beta\in[0,1]} \alpha\beta + (1-\alpha)\cdot(1-0.172\beta)
\end{aligned}$$

$$= 0.853,$$

where the first equality follows by plugging in the values of $\mathrm{OPT}(S_1)$ and $\mathrm{OPT}(S_2)$, and the lone inequality follows via the change of variable $\beta = \frac{k}{m}$. We allow $\beta \in [0,1]$, which relaxes the inner maximization, thus leading to the inequality. Finally, the last equality can be verified with straightforward calculus.

A.2. Proof of Lemma 1

Before moving to the proof of the lemma, we first present an intermediate claim, whose proof appears at the end of this section.

Intermediate claim. The following claim establishes that higher deposits necessarily beget higher rewards.

CLAIM 3. *For shippers t, t' with deposits that satisfy $d_t \geq d_{t'}$ (and hence $p_t \geq p_{t'}$), we must have that $r(d_t, p_t) \geq r(d_{t'}, p_{t'})$.*

Proof of the lemma. We prove the lemma by considering two cases based on the number of arriving reliable shippers. For ease of notation, we suppress all dependencies on S.

- Case 1 - $|T_{\text{rel}}| \geq m$: In this case, $\mathcal{A}^*$ consists of all the unreliable shippers, as well as the m reliable shippers with the m largest rewards, which we have denoted as $T^*_{\text{rel}}(m)$. We will show that the expected profit of $\mathcal{A}^*$ is larger than that of any alternative acceptance set $A \subseteq [T]$. Using (2), we have that

$$\begin{aligned}
\mathbb{E}\left[R(\mathcal{A})\right] &= \sum_{\substack{t\in\mathcal{A}:\\ \rho_{\mathcal{A}}(t)\leq m}} r(d_t,p_t) + \sum_{\substack{t\in\mathcal{A}:\\ \rho_{\mathcal{A}}(t)> m}} \delta(d_t,p_t) \\
&= \sum_{\substack{t\in\mathcal{A}:\\ \rho_{\mathcal{A}}(t)\leq m}} r(d_t,p_t) + \sum_{\substack{t\in\mathcal{A}\cap T_{\text{rel}}:\\ \rho_{\mathcal{A}}(t)> m}} \delta(d_t,p_t) + \sum_{\substack{t\in\mathcal{A}\cap T_{\text{unrel}}:\\ \rho_{\mathcal{A}}(t)> m}} \delta(d_t,p_t) \\
&\leq \sum_{\substack{t\in\mathcal{A}:\\ \rho_{\mathcal{A}}(t)\leq m}} r(d_t,p_t) + \sum_{\substack{t\in\mathcal{A}\cap T_{\text{unrel}}:\\ \rho_{\mathcal{A}}(t)> m}} \delta(d_t,p_t) \\
&\leq \sum_{\substack{t\in\mathcal{A}:\\ \rho_{\mathcal{A}}(t)\leq m}} r(d_t,p_t) + \sum_{\substack{t\in T_{\text{unrel}}:\\ \rho_{\mathcal{A}}(t)> m}} \delta(d_t,p_t) \\
&\leq \sum_{t\in T^*_{\text{rel}}(m)} r(d_t,p_t) + \sum_{\substack{t\in T_{\text{unrel}}:\\ \rho_{\mathcal{A}}(t)> m}} \delta(d_t,p_t) \\
&= \mathbb{E}\left[R(\mathcal{A}^*)\right],
\end{aligned}$$

where the first inequality follows since the penalty of reliable shippers is non-positive, while the second inequality follows because the penalty of unreliable shippers is non-negative. The final inequality uses the fact that $T^*_{\text{rel}}(m)$ is comprised of the m shippers whose rewards are largest, given Claim 3.

- Case 2 - $|T_{\text{rel}}| < m$: In this case, we have that $\mathcal{A}^* = [T]$. Consider any acceptance set $\mathcal{A}$ that includes all unreliable shippers and that satisfies $\mathcal{A} \cap T_{\text{rel}} \subset T_{\text{rel}}$, meaning that not all reliable shippers have been accepted. For any shipper $\tau \in T_{\text{rel}}$ such that $\tau \notin \mathcal{A}$, define the acceptance set $\mathcal{A}' = A \cup \{\tau\}$. We will prove that $\mathbb{E}[R(\mathcal{A}')] \geq \mathbb{E}[R(\mathcal{A})]$. We can then iteratively apply this result until all shippers have been accepted to establish the lemma statement. If $|\mathcal{A}| < m$, then the result trivially holds since

$$\mathbb{E}[R(\mathcal{A}')] - \mathbb{E}[R(\mathcal{A})] = r(d_\tau, p_\tau) \geq 0.$$

Otherwise, let $\ell = \rho^{-1}_{\mathcal{A}}(m)$ denote the unreliable shipper with the m-the largest deposit within $\mathcal{A}$, noting that this shipper must indeed be unreliable, since we have assumed that fewer than m reliable shippers arrive. We have that

$$\begin{aligned}\mathbb{E}[R(\mathcal{A}')] - \mathbb{E}[R(\mathcal{A})] &= r(d_\tau, p_\tau) - r(d_\ell, p_\ell) + \delta(d_\ell, p_\ell) \\ &\geq r(d_\tau, p_\tau) - r(d_\ell, p_\ell) \\ &\geq 0,\end{aligned}$$

where the first inequality uses the fact that the penalty of all unreliable shippers is non-negative. To see the second inequality, observe that since τ is reliable and ℓ is unreliable, it must be the case that $d_\tau \geq d_\ell$. In turn, this allows us to invoke Claim 3 to show that $r(d_\tau, p_\tau) \geq r(d_\ell, p_\ell)$.

Proof of Claim 3. Let $\epsilon_d = d_t - d_{t'}$ and $\epsilon_p = p_t - p_{t'}$, noting that $\epsilon_d, \epsilon_p \geq 0$ since show-up probabilities must be non-decreasing in the deposit. We have that

$$\begin{aligned}r(d_t, p_t) = p_t + (1 - p_t) \cdot d_t &= p_{t'} + \epsilon_p + (1 - p_{t'} - \epsilon_p) \cdot (d_{t'} + \epsilon_d) \\ &= p_{t'} + (1 - p_{t'}) \cdot d_{t'} + \epsilon_p - \epsilon_p d_{t'} - \epsilon_p \epsilon_d \\ &= r(d_{t'}, p_{t'}) + \epsilon_p \cdot (1 - \underbrace{(d_{t'} + \epsilon_d)}_{=d_t \leq 1}) \\ &\geq r(d_{t'}, p_{t'}),\end{aligned}$$

as desired.

A.3. Proof of Claim 1

The worst-case penalty $\delta^*(r)$ can be expressed as the optimal objective of the following optimization problem

$$\delta^*(r) = \min_{p\in[0.5,r]} (1-2p)\cdot(\frac{r-p}{1-p}),$$

which uses the fact that $d = \frac{r-p}{1-p}$. Straightforward calculus can then be used to show that the optimal solution to the above problem matches the structure of the solution proposed in the claim statement.

A.4. Proof of Lemma 2

Let $\mathcal{U}^+ = \{(t,k,\mathcal{A},\theta) : t\in[T], k\in\{1,2,3,4\}, \mathcal{A}\in[0.5,1]^k, \theta\in\{\mathrm{OFF}\}\cup[\max\mathcal{A},1]\}$ denote our original universe of states augmented by a shipper index $t\in[T]$. Moreover, for our fixed reward tuple sequence S, let $\mathcal{U}^+(S)\subset\mathcal{U}^+$ be the set of all states that we visit with a non-zero probability when rolling out our RT policy π on sequence S. To prove the lemma, we will show that for any $(t,k,\mathcal{A},\theta)\in\mathcal{U}^+(S)$, we have that

$$\frac{J^S_\pi(t,k,\mathcal{A},\theta)}{\mathrm{OPT}(S)} \geq V(k,\mathcal{A},\theta). \tag{5}$$

Applying this result for the initial state $(2,1,\{r_1\},\mathrm{OFF})$, we get that

$$\mathrm{CR}_\pi(S) = \frac{\mathbb{E}[R(\mathcal{A}_\pi(S))]}{\mathrm{OPT}(S)} = \frac{J^S_\pi(2,1,\{r_1\},\mathrm{OFF})}{\mathrm{OPT}(S)} \geq V(1,\{r_1\},\mathrm{OFF}) \geq \mathrm{CR}_{\mathrm{RT}},$$

as desired. We begin with an intermediate claim that will prove critical in the induction argument to establish (5).

Intermediate claims. The following intermediate claim, whose proofs can be found at the end of this section, establishes that each threshold serves as an upper bound on the highest reward shipper seen thus far.

CLAIM 4. *For any $(t,k,\mathcal{A},\theta)\in\mathcal{U}^+(S)$, we have that*

$$\max_{\tau\in[t-1]} r_\tau \leq \begin{cases} \theta & \text{if } \theta\neq\mathrm{OFF} \\ \theta_{\mathcal{A}} & \text{if } \theta=\mathrm{OFF}. \end{cases}$$

Proof of the lemma. We proceed to showing (5) via induction over t, starting with the base case of $t=T+1$, where we consider two cases based on whether or not $\theta=\mathrm{OFF}$.

- Case 1 - $\theta=\mathrm{OFF}$: In this case, looking at (3), we have that

$$V(k,\mathcal{A},\mathrm{OFF}) \leq \frac{\mathbb{E}[R(\mathcal{A})]}{\theta_{\mathcal{A}}}$$

$$= \frac{J_\pi^S(T+1,k,\mathcal{A},\mathrm{OFF})}{\theta_{\mathcal{A}}}$$
$$\leq \frac{J_\pi^S(T+1,k,\mathcal{A},\mathrm{OFF})}{\mathrm{OPT}(S)},$$

where the last inequality uses Claim 4, along with the fact that $\mathrm{OPT}(S) = \max_{\tau\in[T]} r_\tau$.

- Case 2 - $\theta \neq \mathrm{OFF}$: Here, exploiting (4) and Claim 4, we have that

$$V(k,\mathcal{A},\theta) \leq \frac{\mathbb{E}\left[R(\mathcal{A})\right]}{\theta}$$
$$= \frac{J_\pi^S(T+1,k,\mathcal{A},\theta)}{\theta}$$
$$\leq \frac{J_\pi^S(T+1,k,\mathcal{A},\theta)}{\mathrm{OPT}(S)},$$

as desired. We now move to the general case of $t < T+1$, where there will once again be two cases based on whether or not $\theta = \mathrm{OFF}$. Each of these two cases with be broken down into relevant sub-cases that dictate how the roll-out policy is executed for shipper t.

- Case 1a - $\theta = \mathrm{OFF}$, $\theta_{\mathcal{A}} = \infty$: When $\theta_{\mathcal{A}} = \infty$, it is optimal to stop accepting any future shippers and so we have that

$$V(k,\mathcal{A},\mathrm{OFF}) = \mathbb{E}\left[R(\mathcal{A})\right] = J_\pi^S(t,k,\mathcal{A},\mathrm{OFF}),$$

and hence (5) holds, since $\mathrm{OPT}(S) \leq 1$.

- Case 1b - $\theta = \mathrm{OFF}$, $r_t \leq \theta_{\mathcal{A}} \in [0.5, 1]$: In this case, we reject shipper t, so we have that

$$J_\pi^S(t,k,\mathcal{A},\mathrm{OFF}) = J_\pi^S(t+1,k,\mathcal{A},\mathrm{OFF}) \geq V(k,\mathcal{A},\mathrm{OFF})\cdot \mathrm{OPT}(S),$$

where the inequality uses the induction hypothesis.

- Case 1c - $\theta = \mathrm{OFF}$, $r_t \in [\theta_{\mathcal{A}}, \theta'_{\mathcal{A}}]$: Here, our RT policy dictates that we flip an $\alpha_{\mathcal{A}}$-biased coin to determine if we accept shipper t. As such, we have that

$$\begin{aligned}
&J_\pi^S(t,k,\mathcal{A},\mathrm{OFF})\\
&\quad= \alpha_{\mathcal{A}}\cdot J_\pi^S(t+1,k+1,\mathcal{A}\cup\{r_t\},\mathrm{OFF}) + (1-\alpha_{\mathcal{A}})\cdot J_\pi^S(t+1,k,\mathcal{A},\theta'_{\mathcal{A}})\\
&\quad\geq \left(\alpha_{\mathcal{A}}\cdot V(k+1,\mathcal{A}\cup\{r_t\},\mathrm{OFF}) + (1-\alpha_{\mathcal{A}})\cdot V(k,\mathcal{A},\theta'_{\mathcal{A}})\right)\cdot \mathrm{OPT}(S)\\
&\quad\geq \left(\alpha_{\mathcal{A}}\cdot \min_{r\in[\theta_{\mathcal{A}},\theta'_{\mathcal{A}}]} V(k+1,\mathcal{A}\cup\{r\},\mathrm{OFF}) + (1-\alpha_{\mathcal{A}})\cdot V(k,\mathcal{A},\theta'_{\mathcal{A}})\right)\cdot \mathrm{OPT}(S)\\
&\quad\geq V(k.\mathcal{A},\mathrm{OFF})\cdot \mathrm{OPT}(S),
\end{aligned}$$

where the first inequality follows via the inductive hypothesis, and the second inequality follows since in (3), the outer maximization must be satisfied by one of Cases (ii)-(iv), since $\theta_{\mathcal{A}} \neq \infty$.

- Case 1d - $\theta = \text{OFF}$, $r_t \geq \theta'_{\mathcal{A}}$: In this final case, shipper t is accepted with certainty, and thus

$$\begin{aligned}
J_\pi^S(t,k,\mathcal{A},\text{OFF}) &= J_\pi^S(t+1,k+1,\mathcal{A}\cup\{r_t\},\text{OFF}) \\
&\geq V(k+1,\mathcal{A}\cup\{r_t\},\text{OFF})\cdot\text{OPT}(S) \\
&\geq \min_{r\geq[\theta'_{\mathcal{A}},1]} V(k+1,\mathcal{A}\cup\{r\},\text{OFF})\cdot\text{OPT}(S) \\
&\geq V(k,\mathcal{A},\text{OFF})\cdot\text{OPT}(S).
\end{aligned}$$

- Case 2a - $\theta \in [0.5,1]$, $r_t \leq \theta$: In this case, we reject shipper t, so we have that

$$J_\pi^S(t,k,\mathcal{A},\theta) = J_\pi^S(t+1,k,\mathcal{A},\theta) \geq V(k,\mathcal{A},\theta)\cdot\text{OPT}(S),$$

where the inequality uses the induction hypothesis.

- Case 2b - $\theta \in [0.5,1]$, $r_t > \theta$: Here, we accept shipper t with certainty and so

$$\begin{aligned}
J_\pi^S(t,k,\mathcal{A},\theta) &= J_\pi^S(t+1,k+1,\mathcal{A}\cup\{r_t\},\text{OFF}) \\
&\geq V(k+1,\mathcal{A}\cup\{r_t\},\text{OFF})\cdot\text{OPT}(S) \\
&\geq \min_{r\geq[\theta,1]} V(k+1,\mathcal{A}\cup\{r\},\theta)\cdot\text{OPT}(S) \\
&\geq V(k,\mathcal{A},,\theta)\cdot\text{OPT}(S),
\end{aligned}$$

where the first inequality uses the induction hypothesis and the last inequality uses (4).

Proof of Claim 4 To prove the claim, we consider an arbitrary realization of the sequence of states visited under our RT policy π

$$(2,\underbrace{k^{(2)}}_{=1},\underbrace{\mathcal{A}^{(2)}}_{=\{r_1\}},\underbrace{\theta^{(2)}}_{=\text{OFF}}) \quad\rightarrow\quad (3,k^{(3)},\mathcal{A}^{(3)},\theta^{(3)}) \quad\rightarrow\quad \cdots \quad\rightarrow\quad (T+1,k^{(T+1)},\mathcal{A}^{(T+1)},\theta^{(T+1)}),$$

and we prove the claim statement for each of the realized states. We proceed by induction over t with a base case of $t = 2$. Since $\theta^{(2)} = \text{OFF}$, we need to show that $\theta_{\mathcal{A}^{(2)}} \geq r_1$, which holds trivially, since $\mathcal{A}^{(2)} = \{r_1\}$ and so $\theta_{\mathcal{A}^{(2)}} \in [r_1, 1]$. We now proceed to the general case of $t > 2$ by considering four cases based on the value of r_t relative to the shipper $t-1$ thresholds.

- Case 1 - $\theta^{(t-1)} = \text{OFF},\ r_{t-1} < \theta_{\mathcal{A}^{(t-1)}}$: In this case , shipper $t-1$ is rejected and thus we know that $\mathcal{A}^{(t-1)} = \mathcal{A}^{(t)}$ and that $\theta^{(t-1)} = \theta^{(t)} = \text{OFF}$. Using the induction hypothesis along with the fact that the thresholds depend only on the current acceptance set, we have that

$$\theta_{\mathcal{A}^{(t)}} = \theta_{\mathcal{A}^{(t-1)}} \geq \max\{r_{t-1}, \max_{\tau\in[t-2]} r_\tau\} = \max_{\tau\in[t-1]} r_\tau,$$

as desired since $\theta^{(t)} = \text{OFF}$ as noted above.

- Case 2 - $\theta^{(t-1)} = \text{OFF},\ r_{t-1} \geq \theta_{\mathcal{A}^{(t-1)}}$: If shipper $t-1$ is accepted (either via the randomization or because $r_{t-1} \geq \theta'_{\mathcal{A}^{(t-1)}}$, then $\mathcal{A}^{(t)} = \mathcal{A}^{(t-1)} \cup \{r_{t-1}\}$, and so

$$\begin{aligned}\theta_{\mathcal{A}^{(t)}} \geq \max \mathcal{A}^{(t)} &\geq \max\{r_{t-1}, \theta_{\mathcal{A}^{(t-1)}}\}\\ &\geq \max\{r_{t-1}, \max_{\tau\in[t-2]} r_\tau\}\\ &= \max_{\tau\in[t-1]} r_\tau,\end{aligned}$$

as desired, since $\theta^{(t)} = \text{OFF}$. The second inequality above uses the induction hypothesis. If, on the other hand, shipper t is rejected, it must have been because the $\alpha_{\mathcal{A}^{(t-1)}}$-biased coin flip suggested so. The result of such a rejection is that $\theta^{(t)} = \theta'_{\mathcal{A}^{(t-1)}}$, i.e., we move to shipper t with an active threshold of $\theta'_{\mathcal{A}^{(t-1)}}$. As a result, we have that

$$\begin{aligned}\theta^{(t)} = \theta'_{\mathcal{A}^{(t-1)}} &\geq \max\{r_{t-1}, \theta_{\mathcal{A}^{(t-1)}}\}\\ &\geq \max\{r_{t-1}, \max_{\tau\in[t-2]} r_\tau\}\\ &= \max_{\tau\in[t-1]} r_\tau,\end{aligned}$$

where the first inequality uses the fact that $\theta'_{\mathcal{A}^{(t-1)}} \geq r_{t-1}$ since we made rejection only if $r_{t-1} \leq \theta'_{\mathcal{A}^{(t-1)}}$, and the second inequality uses the induction hypothesis.

- Case 3 - $\theta^{(t-1)} \neq \text{OFF},\ r_{t-1} < \theta^{(t-1)}$: In this case, shipper $t-1$ is rejected, $\theta^{(t-1)} = \theta^{(t)}$ and so,

$$\theta^{(t)} = \theta^{(t-1)} \geq \max\{r_{t-1}, \max_{\tau\in[t-2]} r_\tau\} = \max_{\tau\in[t-1]} r_\tau,$$

where the inequality uses the induction hypothesis.

- Case 4 - $\theta^{(t-1)} \neq \text{OFF},\ r_{t-1} \geq \theta^{(t-1)}$: In this final case, shipper $t-1$ is accepted with certainty and so $\mathcal{A}^{(t)} = \mathcal{A}^{(t-1)} \cup \{r_{t-1}\}$ and $\theta^{(t)} = \text{OFF}$. Exploiting these facts, we see that

$$\begin{aligned}\theta_{\mathcal{A}^{(t)}} \geq \max \mathcal{A}^{(t)} &\geq \max\{r_{t-1}, \theta^{(t-1)}\}\\ &\geq \max\{r_{t-1}, \max_{\tau\in[t-2]} r_\tau\}\\ &= \max_{\tau\in[t-1]} r_\tau,\end{aligned}$$

where the second inequality uses the induction hypothesis.

A.5. Proof of Lemma 3

We prove the result via induction over k. For the base case of $kh = 4$, we have that

$$V(k,\mathcal{A},\cdot) = \mathbb{E}\left[R(\mathcal{A})\right] \geq \mathbb{E}\left[\hat{R}(\mathcal{A})\right] = \mathbb{E}\left[\hat{R}(G_\epsilon(\mathcal{A}))\right] = \hat{V}(k, G_\epsilon(\mathcal{A}),\cdot),$$

as desired. Moving the general case of $k \in \{1,2,3\}$, we consider two cases based on whether or not θ is an active threshold.

- Case 1 - $\theta \neq \text{OFF}$: We have that

$$\begin{aligned} V(k,\mathcal{A},\theta) &= \min\left\{\frac{\mathbb{E}\left[R(\mathcal{A})\right]}{\theta}, \min_{r\in[\theta,1]} V(k+1, A\cup\{r\},\text{OFF})\right\} \\ &\geq \min\left\{\frac{\mathbb{E}\left[\hat{R}(G_\epsilon(\mathcal{A}))\right]}{G_\epsilon(\theta)}, \min_{r\in[\theta,1]} \hat{V}(k+1, G_\epsilon(A\cup\{r\}),\text{OFF})\right\} \\ &= \hat{V}\left(k, G_\epsilon(\mathcal{A}), G_\epsilon(\theta)\right), \end{aligned}$$

  where the inequality uses the induction hypothesis.

- Case 2 - $\theta = \text{OFF}$: In this second case, we have that

$$\begin{aligned} V(k,\mathcal{A},\text{OFF}) \geq \max\Bigg\{&\mathbb{E}\left[R(\mathcal{A})\right], \max_{\substack{\max\mathcal{A}\leq\theta_\mathcal{A}\leq\theta'_\mathcal{A}\leq 1,\\ \alpha_\mathcal{A}\in G_\epsilon}} \min\Bigg\{\frac{\mathbb{E}\left[R(\mathcal{A})\right]}{\theta_\mathcal{A}}, \\ &\min_{r\in[\theta_\mathcal{A},\theta'_\mathcal{A}]} \alpha_\mathcal{A}\cdot V(k+1,\mathcal{A}\cup\{r\},\text{OFF}) + (1-\alpha_\mathcal{A})\cdot V(k,\mathcal{A},\theta'_\mathcal{A}), \\ &\min_{r\in[\theta'_\mathcal{A},1]} V(k+1,\mathcal{A}\cup\{r\},\text{OFF})\Bigg\}\Bigg\}, \end{aligned}$$

  where the inequality follows by restricting $\alpha_\mathcal{A}$ to the grid G_ϵ. Additionally, let $\hat{\theta}_\mathcal{A}, \hat{\theta}'_\mathcal{A}, \hat{\alpha}_\mathcal{A}$ denote the optimal thresholds and randomization probability for the discretized DP at the state $(k, G_\epsilon(\mathcal{A}), \text{OFF})$, which is feasible above, since $\max \mathcal{G}_\epsilon(\mathcal{A}) \geq \max\mathcal{A}$. Accordingly, we have that

$$\begin{aligned} &V(k,\mathcal{A},\text{OFF}) \\ &\geq \max\Bigg\{\mathbb{E}\left[R(\mathcal{A})\right], \min\Bigg\{\frac{\mathbb{E}\left[R(\mathcal{A})\right]}{\hat{\theta}_\mathcal{A}}, \\ &\quad\min_{r\in[\hat{\theta}_\mathcal{A},\hat{\theta}'_\mathcal{A}]} \hat{\alpha}_\mathcal{A}\cdot V(k+1,\mathcal{A}\cup\{r\},\text{OFF}) + (1-\hat{\alpha}_\mathcal{A})\cdot V(k,\mathcal{A},\hat{\theta}'_\mathcal{A}), \\ &\quad\min_{r\in[\hat{\theta}'_\mathcal{A},1]} V(k+1,\mathcal{A}\cup\{r\},\text{OFF})\Bigg\}\Bigg\} \end{aligned}$$

$$
\begin{aligned}
&\geq \max\left\{\mathbb{E}\left[\hat{R}(G_\epsilon(\mathcal{A}))\right], \min\left\{\frac{\mathbb{E}\left[\hat{R}(G_\epsilon(\mathcal{A}))\right]}{G_\epsilon(\hat{\theta}_{\mathcal{A}})},\right.\right.\\
&\quad \min_{r\in[\hat{\theta}_{\mathcal{A}},\hat{\theta}'_{\mathcal{A}}]} \hat{\alpha}_{\mathcal{A}}\cdot\hat{V}\left(k+1, G_\epsilon(\mathcal{A}\cup\{r\}), \mathrm{OFF}\right) + (1-\hat{\alpha}_{\mathcal{A}})\cdot\hat{V}\left(k, G_\epsilon(\mathcal{A}), G_\epsilon(\hat{\theta}'_{\mathcal{A}})\right),\\
&\quad \left.\left.\min_{r\in[\hat{\theta}'_{\mathcal{A}},1]} \hat{V}(k+1, G_\epsilon(\mathcal{A}\cup\{r\}), \mathrm{OFF})\right\}\right\}\\
&= \hat{V}\left(k, G_\epsilon(\mathcal{A}), \mathrm{OFF}\right),
\end{aligned}
$$

where the second inequality uses the induction hypothesis as well as the result from Case 1 for the active threshold.

A.6. Proof of Lemma 4

For any reward sequence S, the competitive ratio of $\pi_{(m)}$ is

$$
\begin{aligned}
\frac{\mathbb{E}\left[R\left(\mathcal{A}_{\pi_{(m)}}(S)\right)\right]}{\mathrm{OPT}(S)} &\geq \frac{\mathbb{E}\left[\sum_{k\in[m]} R\left(A_{\pi_{(1)}}(S_k)\right)\right]}{\mathrm{OPT}(S)}\\
&= \sum_{k\in[m]} \frac{\mathbb{E}\left[R\left(A_{\pi_{(1)}}(S_k)\right)\right]}{r_{\max}(k)}\cdot\frac{r_{\max}(k)}{\mathrm{OPT}(S)}\\
&\geq \sum_{k\in[m]} 0.819\cdot\frac{r_{\max}(k)}{\mathrm{OPT}(S)}\\
&= \sum_{k\in[m]} 0.819\cdot\frac{r_{\max}(k)}{\sum_{k\in[m]} r_{\max}(k)}\\
&= 0.819,
\end{aligned}
$$

where the first inequality follows since the DO slot allocation mechanism will only improve our expected profit, while the second inequality employs Theorem 1. The second equality uses the fact that the clairvoyant's expected profit can be computed by summing the highest reward shipper assigned to each slot.

Appendix B: Missing Proofs from Section 4

B.1. Proof of Claim 2

We show first that $\mathrm{OPT}_{\text{fluid}} \geq \mathrm{OPT}$. For this purpose, let $\mathcal{A}^* = \arg\max_{\mathcal{A}\subseteq[T]} \mathcal{R}(\mathcal{A}; m)$ denote the optimal clairvoyant acceptance set, and note that

$$\mathrm{OPT} = \sum_{t \in \mathcal{A}^*} \left(\underbrace{\Pr\left[\sum_{\tau:\rho(\tau)<\rho(t)} Y_\tau < m \right] \cdot p_t}_{\text{service fee}} - \underbrace{\Pr\left[\sum_{\tau:\rho(\tau)<\rho(t)} Y_\tau \geq m \right] d_t p_t}_{\text{rolled/overbooked}} + \underbrace{(1-p_t)\cdot d_t}_{\text{no-show}} \right).$$

To prove this first bound, we will define a feasible solution $(\hat{x}, \hat{y}, \hat{z})$ to Fluid IP that achieves an objective value of OPT. To do so, for $t \notin \mathcal{A}^*$, we set $\hat{x}_t = \hat{y}_t = \hat{z}_t = 0$, and for $t \in \mathcal{A}^*$, we set

$$\begin{aligned}
\hat{x}_t &= 1 \\
\hat{y}_t &= \Pr\left[\sum_{\tau:\rho(\tau)<\rho(t)} Y_\tau < m \right] \cdot p_t \\
\hat{z}_t &= \Pr\left[\sum_{\tau:\rho(\tau)<\rho(t)} Y_\tau \geq m \right] \cdot p_t.
\end{aligned}$$

It is straightforward to see that this solution yields an objective of OPT, and that constraint (2) holds. To see that constraint (1) holds, observe that

$$\begin{aligned}
\sum_{t\in[T]} \hat{y}_t &= \sum_{t\in\mathcal{A}^*} \Pr\left[\sum_{\tau:\rho(\tau)<\rho(t)} Y_\tau < m \right] \cdot p_t \\
&= \mathbb{E}\left[\min\left\{ \sum_{t\in\mathcal{A}^*} Y_t, m \right\} \right] \\
&\leq m.
\end{aligned}$$

Next, we prove the second bound, namely that $2m + \sum_{t\in T_{\text{unrel}}} d_t \cdot (1-2p_t) \geq \mathrm{OPT}_{\text{fluid}}$. Let (x^*, y^*, z^*) denote the optimal solution to Fluid IP, and observe that

$$\begin{aligned}
\mathrm{OPT}_{\text{fluid}} &= \sum_{t:x_t^*=1} y_t^* + \sum_{t:x_t^*=1} d_t \cdot \left((1-p_t) - (p_t - y_t^*)\right) \\
&= \sum_{t:x_t^*=1} y_t^* + \sum_{t:x_t^*=1} d_t \cdot (1-2p_t) + \sum_{t:x_t^*=1} d_t y_t^* \\
&\leq 2\sum_{t\in[T]} y_t + \sum_{t:x_t^*=1} d_t \cdot (1-2p_t) \\
&\leq 2\sum_{t\in[T]} y_t + \sum_{t\in T_{\text{unrel}}} d_t \cdot (1-2p_t) \\
&\leq 2\min\{\sum_{t\in[T]} p_t, m\} + \sum_{t\in T_{\text{unrel}}} d_t \cdot (1-2p_t),
\end{aligned}$$

where the first inequality uses the fact that $d_t \leq 1$, and the second inequality follows because $1 - 2p_t \geq 0$ if and only if $t \in T_{\text{unrel}}$. The last inequality follows since constraint (1) implies that $\sum_{t\in[T]} y_t \leq m$, and constraint (2) implies that $y_t \leq p_t$ and hence $\sum_{t\in[T]} y_t \leq \sum_{t\in[T]} p_t$.

B.2. Proof of Theorem 2

The reliable first mechanism. We will analyze the performance of the GUM policy using the reliable first (slot) allocation mechanism. At a high level, the RF mechanism works exactly like the DO mechanism except that it prioritizes all reliable shippers before those that are unreliable. Formally, we define this alternative slot allocation mechanism works as follows.

- The slots are first allocated to the reliable shippers who show up in decreasing order of their deposit size.
- The remaining slots are then allocated to the unreliable shippers who show up in decreasing order of deposit size.

We reiterate that we use this RF mechanism solely for the purpose of our analysis. In practice, the liner will actually implement the optimal DO slot allocation mechanism. We will use $\hat{\mathcal{R}}(\mathcal{A})$ to denote the expected profit earned from acceptance set $\mathcal{A}$ under the RF mechanism.

Intermediate results. Our proof critically relies on the following three results, all of which concern properties of binomial random variables. Throughout this section, we use the shorthand Binom(n, p) to represent a binomial random variable with n trials, each with success probability p.

CLAIM 5. *Let $X = \sum_{t\in[T]} B_t$, where $B_t \sim \text{Bernoulli}(p_t)$ are independent Bernoulli trials with success probabilities $\sum_{t\in[T]} p_t = \mu$. For arbitrary integer $m \geq 0$, we have*

$$\mathbb{E}\left[\min\{X, m\}\right] \geq \mathbb{E}\left[\min\left\{\text{Binom}\left(\frac{\mu}{T}, T\right), m\right\}\right].$$

CLAIM 6 (**Ma et al. (2021) Lemma EC.5**). *Let c be any real-valued positive number, T any positive integer. The function*

$$f(\lambda) = \frac{\mathbb{E}\left[\min\{c, \text{Binom}(T, \lambda/T)\}\right]}{\lambda}$$

is non-increasing in λ over the interval $[0, T]$.

CLAIM 7 (**Alaei et al. (2012) Lemma 5.3**). *For positive integer T and $\lambda \in [0, T]$, let $X \sim \text{Binom}(T, \frac{\lambda}{T})$, then $\mathbb{E}[\min\{X, b\}] \geq \text{APX}(b) \cdot \lambda$ for any $b \geq \lambda$.*

Proof of theorem. Let $\hat{R}_{\mathrm{rel}}$ and $\hat{R}_{\mathrm{unrel}}$ respectively denote the expected profit accrued by reliable and unreliable shippers under the GUM policy, and when the slots are allocated sub-optimally according to the RF mechanism. Accordingly, we have that $\hat{\mathcal{R}}(\mathcal{A}_{\mathrm{GUM}}) = \hat{R}_{\mathrm{rel}} + \hat{R}_{\mathrm{unrel}}$. Our proof initially proceeds by establishing lower bounds on the expected profit contribution of the two shipper types.

Bounding reliable profit. Under the RF mechanism, we have that

$$
\begin{aligned}
\hat{R}_{\mathrm{rel}} &= \sum_{t \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}} \left(d_t \cdot (1 - p_t) + \Pr\left[\sum_{\substack{\tau \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}: \\ \rho(\tau) < \rho(t)}} Y_\tau < m \right] \cdot p_t \right. \\
&\qquad \left. - \Pr\left[\sum_{\substack{\tau \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}: \\ \rho(\tau) < \rho(t)}} Y_\tau \geq m \right] \cdot p_t d_t \right) \\
&= \sum_{t \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}} \left(\Pr\left[\sum_{\substack{\tau \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}: \\ \rho(\tau) < \rho(t)}} Y_\tau < m \right] \cdot p_t \right. \\
&\qquad \left. + \sum_{t \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}} d_t \cdot (1 - p_t) - \sum_{\substack{\tau \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}: \\ \rho(\tau) > m}} p_t d_t \right) \\
&\geq \sum_{t \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}} \left(\Pr\left[\sum_{\substack{\tau \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}: \\ \rho(\tau) < \rho(t)}} Y_\tau < m \right] \cdot p_t \right. \\
&\qquad \left. + \sum_{\substack{\tau \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}: \\ \rho(\tau) \leq m}} d_t \cdot (1 - p_t) - \sum_{\substack{\tau \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}: \\ \rho(\tau) > m}} p_t d_t \right) \\
&\geq \underbrace{\sum_{t \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}} \Pr\left[\sum_{\substack{\tau \in \mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}: \\ \rho(\tau) < \rho(t)}} Y_\tau < m \right] \cdot p_t}_{\text{Bound 1}}
\end{aligned}
$$

The second equality results because any shipper $t \in [T]$ with deposit rank $\rho(t) \leq m$ cannot be overbooked. The first inequality results ignoring the no-show profit of reliable shippers whose rank is larger than m. The last inequality is trivial if $|\mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}| \leq m$, since in this case, there will not be no overbooking penalty. To the see this last inequality when we accept strictly more than

m reliable shippers, define $\rho^{-1}(m)$ to be the m-th largest deposit among $\mathcal{A}_{\text{GUM}} \cap T_{\text{rel}}$, and observe that

$$\begin{aligned}
\sum_{\substack{\tau \in \mathcal{A}_{\text{GUM}} \cap T_{\text{rel}}: \\ \rho(\tau) \leq m}} d_t \cdot (1 - p_t) - \sum_{\substack{\tau \in \mathcal{A}_{\text{GUM}} \cap T_{\text{rel}}: \\ \rho(\tau) > m}} p_t d_t &\geq d_{\rho^{-1}(m)} \cdot \left(\sum_{\substack{\tau \in \mathcal{A}_{\text{GUM}} \cap T_{\text{rel}}: \\ \rho(\tau) \leq m}} (1 - p_t) - \sum_{\substack{\tau \in \mathcal{A}_{\text{GUM}} \cap T_{\text{rel}}: \\ \rho(\tau) > m}} p_t \right) \\
&= d_{\rho^{-1}(m)} \cdot \left(m - \sum_{t \in \mathcal{A}_{\text{GUM}}} p_t \right) \\
&\geq 0,
\end{aligned}$$

where the first inequality follows since for any reliable shipper $t \in \mathcal{A}_{\text{GUM}} \cap T_{\text{rel}}$ such that $\rho(t) \leq m$, we have that $d_t \geq d_{\rho^{-1}(m)}$, while the opposite inequality is true with respect to the deposits of shippers with rank higher strictly larger than m. The last inequality follows by observing that the GUM policy never accepts more than m units of fluid demand.

Bounding unreliable profit. Initially using the fact that the GUM policy accepts all unreliable shippers, we have that

$$\begin{aligned}
\hat{R}_{\text{unrel}} &= \sum_{t \in T_{\text{unrel}}} \left(d_t \cdot (1 - 2p_t) + \Pr\left[\sum_{\tau \in \mathcal{A} \cap T_{\text{rel}}} Y_\tau + \sum_{\substack{\tau \in \mathcal{A} \cap T_{\text{unrel}}: \\ \rho(\tau) < \rho(t)}} Y_\tau < m \right] \cdot (p_t \cdot (1 + d_t)) \right) \\
&\geq \underbrace{\sum_{t \in T_{\text{unrel}}} d_t \cdot (1 - 2p_t) + \sum_{t \in T_{\text{unrel}}} \Pr\left[\sum_{\tau \in \mathcal{A} \cap T_{\text{rel}}} Y_\tau + \sum_{\substack{\tau \in \mathcal{A} \cap T_{\text{unrel}}: \\ \rho(\tau) < \rho(t)}} Y_\tau < m \right] \cdot p_t}_{\text{Bound 2}},
\end{aligned}$$

where the lone inequality follows because $d_t \geq 0$.

Establishing the competitive ratio. Combining the two bounds, we get that

$$\begin{aligned}
\mathcal{R}(\mathcal{A}_{\text{GUM}}) &\geq \sum_{t \in T_{\text{unrel}}} d_t \cdot (1 - 2p_t) + \sum_{t \in T_{\text{unrel}}} \Pr\left[\sum_{\tau \in \mathcal{A} \cap T_{\text{rel}}} Y_\tau + \sum_{\substack{\tau \in \mathcal{A} \cap T_{\text{unrel}}: \\ \rho(\tau) < \rho(t)}} Y_\tau < m \right] \cdot p_t \\
&\quad + \sum_{t \in \mathcal{A}_{\text{GUM}} \cap T_{\text{rel}}} \Pr\left[\sum_{\substack{\tau \in \mathcal{A}_{\text{GUM}} \cap T_{\text{rel}}: \\ \rho(\tau) < \rho(t)}} Y_\tau < m \right] \cdot p_t \\
&= \sum_{t \in T_{\text{unrel}}} d_t \cdot (1 - 2p_t) + \sum_{t \in \mathcal{A}_{\text{GUM}}} \Pr\left[\sum_{\substack{\tau \in \mathcal{A}_{\text{GUM}}: \\ \rho(\tau) < \rho(t)}} Y_\tau < m \right] \cdot p_t
\end{aligned}$$

$$= \sum_{t \in T_{\mathrm{unrel}}} d_t \cdot (1 - 2p_t) + \mathbb{E}\left[\min\left\{\sum_{t \in \mathcal{A}_{\mathrm{GUM}}} Y_t, m\right\}\right],$$

where the first equality follows because $\mathcal{A}_{\mathrm{GUM}} = (\mathcal{A}_{\mathrm{GUM}} \cap T_{\mathrm{rel}}) \cup T_{\mathrm{unrel}}$ combined with the fact that the latter two terms are precisely the expected number of slots allocated, which is agnostic to the allocation mechanism, i.e. for any acceptance set, the expected number of slots allocated is the same under the DO and RF mechanisms. Finally, letting $\mu = \sum_{t \in \mathcal{A}_{\mathrm{GUM}}} p_t$, and invoking Claim 5, we continue this string of inequalities to yield

$$\mathcal{R}(\mathcal{A}_{\mathrm{GUM}}; m) \geq \sum_{t \in T_{\mathrm{unrel}}} d_t \cdot (1 - 2p_t) + \mathbb{E}\left[\min\left\{\mathrm{Binom}\left(|\mathcal{A}_{\mathrm{GUM}}|, \frac{\mu}{|\mathcal{A}_{\mathrm{GUM}}|}\right), m\right\}\right].$$

In turn, using the upper bound on OPT given in Claim 2 yields a competitive ratio of

$$\begin{aligned}
\frac{\mathcal{R}(\mathcal{A}_{\mathrm{GUM}})}{\mathrm{OPT}} &\geq \frac{1}{2} \cdot \frac{\sum_{t \in T_{\mathrm{unrel}}} d_t \cdot (1 - 2p_t) + \mathbb{E}\left[\min\left\{\mathrm{Binom}\left(|\mathcal{A}_{\mathrm{GUM}}|, \frac{\mu}{|\mathcal{A}_{\mathrm{GUM}}|}\right), m\right\}\right]}{\sum_{t \in T_{\mathrm{unrel}}} d_t \cdot (1 - 2p_t) + \min\{\sum_{t \in [T]} p_t, m\}} \\
&\geq \frac{1}{2} \cdot \frac{\mathbb{E}\left[\min\left\{\mathrm{Binom}\left(|\mathcal{A}_{\mathrm{GUM}}|, \frac{\mu}{|\mathcal{A}_{\mathrm{GUM}}|}\right), m\right\}\right]}{\min\{\sum_{t \in [T]} p_t, m\}} \\
&\geq \frac{1}{2m} \cdot \mathbb{E}\left[\min\left\{\mathrm{Binom}\left(|\mathcal{A}_{\mathrm{GUM}}|, \frac{\mu}{|\mathcal{A}_{\mathrm{GUM}}|}\right), m\right\}\right] \\
&\geq \frac{1}{2m} \cdot \mathbb{E}\left[\min\left\{\mathrm{Binom}\left(|\mathcal{A}_{\mathrm{GUM}}|, \frac{m-1}{|\mathcal{A}_{\mathrm{GUM}}|}\right), m-1\right\}\right] \\
&= \frac{1}{2} \cdot \frac{m-1}{m} \frac{\mathbb{E}\left[\min\left\{\mathrm{Binom}\left(|\mathcal{A}_{\mathrm{GUM}}|, \frac{m-1}{|\mathcal{A}_{\mathrm{GUM}}|}\right), m-1\right\}\right]}{m-1} \\
&\geq \frac{1}{2} \cdot \mathrm{APX}(m-1).
\end{aligned}$$

The second inequality follows since the ratio clearly becomes smaller when no unreliable shippers arrive. The third inequality follows by observing that when less than m units of fluid demand arrive, we have that $\min\{\sum_{t \in [T]} p_t, m\} = \sum_{t \in \mathcal{A}_{\mathrm{GUM}}} p_t = \mu$, and hence invoking Claim 6, the worst case competitive ratio arises when $\sum_{t \in [T]} p_t \geq m$. From here, we arrive at the fourth inequality by observing that, in the worst case $\mu = m - 1$. Finally, the last inequality follows by Claim 7.

B.3. Proof of Claim 5.

The claim is trivially true if $p_1 = p_2 = \ldots = p_T = \frac{\mu}{T}$, and hence we proceed under the assumption that there exists at least one pair $t_1, t_2 \in [T]$ such that $p_{t_1} > \frac{\mu}{T}$ and $p_{t_2} < \frac{\mu}{T}$. Let $\epsilon = \min\{p_{t_1} - \frac{\mu}{T}, \frac{\mu}{T} - p_{t_2}\}$, and define a new collection of Bernoulli random variables $\{\hat{B}_t\}_{t\in[T]}$, where for $t \in [T] \setminus \{t_1, t_2\}$ we set $\hat{B}_t = B_t$, while $\hat{B}_{t_1} = \text{Bernoulli}(p_{t_1} - \epsilon)$ and $\hat{B}_{t_2} = \text{Bernoulli}(p_{t_2} + \epsilon)$. By construction, it must be the case that at least one of $\hat{B}_{t_1}$ and $\hat{B}_{t_1}$ has success probability of $\frac{\mu}{T}$. Consequently, letting $\hat{X} = \sum_{t\in[T]} \hat{B}_t$, if we can show

$$\mathbb{E}\left[\min\{X,m\}\right] \geq \mathbb{E}\left[\min\left\{\hat{X},m\right\}\right] \tag{6}$$

then the proof is complete, since the above reasoning can be applied iteratively until all success probabilities are $\frac{\mu}{T}$.

To show (6), note that

$$\begin{aligned}
&\mathbb{E}\left[\min\{X,m\}\right] - \mathbb{E}\left[\min\left\{\hat{X},m\right\}\right] = \\
&\quad \underbrace{\mathbb{E}\left[\min\{X,m\} - \min\left\{\hat{X},m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \geq m\right]}_{\text{term 1}} \cdot \Pr\left[\sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \geq m\right] + \\
&\quad \underbrace{\mathbb{E}\left[\min\{X,m\} - \min\left\{\hat{X},m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \leq m-2\right]}_{\text{term 2}} \cdot \Pr\left[\sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \leq m-2\right] + \\
&\quad \underbrace{\mathbb{E}\left[\min\{X,m\} - \min\left\{\hat{X},m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t = m-1\right]}_{\text{term 3}} \cdot \Pr\left[\sum_{t\in[T]\setminus\{t_1,t_2\}} B_t = m = 1\right].
\end{aligned}$$

We proceed to show that each of the three terms is non-positive:

- term 1: For the first term, we have that

$$\mathbb{E}\left[\min\{X,m\} - \min\left\{\hat{X},m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \geq m\right] = 0,$$

  since both X and $\hat{X}$ will be larger than m with probability 1.

- term 2: In this case, we have

$$\begin{aligned}
&\mathbb{E}\left[\min\{X,m\} - \min\left\{\hat{X},m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t \leq m-2\right] \\
&\quad = \mathbb{E}\left[X - \hat{X}\right]
\end{aligned}$$

$$
\begin{aligned}
&= \mathbb{E}\left[\sum_{t\in\{t_1,t_2\}} \left(B_t - \hat{B}_t\right)\right] \\
&= (p_{t_1} + p_{t_2}) - (p_{t_1} - \epsilon + p_{t_2} + \epsilon) \\
&= 0.
\end{aligned}
$$

- term 3: Here, we have

$$
\begin{aligned}
&\mathbb{E}\left[\min\{X, m\} - \min\left\{\hat{X}, m\right\} \,\middle|\, \sum_{t\in[T]\setminus\{t_1,t_2\}} B_t = m-1\right] \\
&\quad= \mathbb{E}\left[\min\{B_{t_1} + B_{t_2}, 1\}\right] - \mathbb{E}\left[\min\{\hat{B}_{t_1} + \hat{B}_{t_2}, 1\}\right] \\
&\quad= (p_{t_1} + (1-p_{t_1})\cdot p_{t_2}) - \left(p_{t_1} - \epsilon + \left(1 - p_{t_1} + \epsilon\right)\cdot\left(p_{t_2} + \epsilon\right)\right) \\
&\quad= \epsilon \cdot \underbrace{(p_{t_2} - p_{t_1})}_{\le -\epsilon} + \epsilon^2 \\
&\quad\le 0
\end{aligned}
$$

Appendix C: A Randomized Asymptotic 0.59-Competitive Approach

In this section, we present a randomized algorithm for the independent show-up setting that achieves a competitive ratio of 0.59 as m tends to infinity. The following theorem formalizes the precise competitive ratio attained by this algorithm; its proof is developed over the remainder of the section.

THEOREM 3. *There exists a randomized algorithm that achieves a competitive ratio of*

$$
\mathrm{CR}(m) = 0.59 \cdot \left(1 - \frac{1}{m-1}\right)^2 \left(2\left(1 - \sqrt{\frac{\log m}{m}}\right)\left(1 - m^{-\frac{1}{2}+o(1)}\right) - 1\right),
$$

noting that

$$
\lim_{m\to\infty} \mathrm{CR}(m) = 0.59.
$$

The algorithm we consider models our online booking problem as an instance of the online knapsack problem and applies the algorithm of Chakrabarty et al. (2008), with a slight modification that we describe below. We therefore begin by outlining this algorithm and its associated performance guarantee in the online knapsack setting. We then explain how the algorithm is adapted to our stochastic overbooking environment.

The threshold-based algorithm of Chakrabarty et al. (2008). Consider an instance of the online knapsack problem, where the value and weight of item j are denoted by v_j and w_j, respectively, and the knapsack has capacity W. Moreover, input parameters $L, U \geq 0$ provide lower and upper bounds on the value-to-weight ratio of all items that may arrive, i.e., the values and weights of each item j are guaranteed to satisfy $L \leq \frac{v_j}{w_j} \leq U$. When item j arrives, its value and weight are observed, and an irrevocable decision must be made regarding whether to pack the item. Letting $z_j \in [0,1]$ denote the fraction of the knapsack capacity that has been filled upon the arrival of item j, Chakrabarty et al. (2008) propose packing item j if doing so does not exceed the knapsack's capacity and $\frac{v_j}{w_j} \geq \Psi(z_j)$, where $\Psi(z) = \left(\frac{Ue}{L}\right)^z \left(\frac{L}{e}\right)$. We refer to this simple threshold-based policy as π_{KNAP}, and we summarize its worst-case performance in the theorem below.

THEOREM 4 **(Theorem 2.1 Chakrabarty et al. (2008)).** *The policy π_{KNAP} achieves a competitive ratio of $\left(\frac{1-\epsilon}{\ln(\frac{U}{L})+1}\right)$, where $\epsilon = \frac{\max_j w_j}{W}$.*

It is important to note that the benchmark used to compute the competitive ratio in the above theorem is the maximum-value feasible packing for the offline problem, in which the identities of all items are revealed upfront.

Our randomized adaptation. The policy we adopt treats each arriving shipper as an item in an instance of the online knapsack problem, where the knapsack capacity is $W = m$. The weight of shipper t is taken to be $w_t = p_t$, while its value is given by

$$v_t = \begin{cases} p_t + (1-p_t) \cdot d_t & \text{if } t \in T_{\text{rel}} \\ p_t \cdot (1+d_t) & \text{if } t \in T_{\text{unrel}}. \end{cases}$$

As such, any feasible packing of shippers $\mathcal{A}$ yields a total value of

$$\begin{aligned} \mathcal{V}(\mathcal{A}) &= \sum_{t \in \mathcal{A} \cap T_{\text{rel}}} \left(p_t + (1-p_t) \cdot d_t\right) + \sum_{t \in \mathcal{A} \cap T_{\text{unrel}}} p_t \cdot (1+d_t) \\ &= \sum_{t \in \mathcal{A}} p_t \cdot (1+d_t) + \sum_{t \in \mathcal{A} \cap T_{\text{rel}}} (1-2p_t) \cdot d_t. \end{aligned}$$

We execute π_{KNAP} on the sequence of arriving shippers using the weights, values, and knapsack capacity described above, and let $\mathcal{A}_{\text{KNAP}}$ denote the subset of "packed" shippers selected by this algorithm. Our final accept/reject decisions for each shipper are as follows:

- We accept all unreliable shippers.

- If π_{KNAP} packs shipper t, i.e., $t \in \mathcal{A}_{\text{KNAP}}$, then we sample $Z_t \sim \text{Bernoulli}(1-\delta)$, where the exact value of δ will be specified shortly to optimize the competitive ratio. If shipper t is reliable, we accept shipper t if $Z_t = 1$ and reject shipper t if $Z_t = 0$. If shipper t is unreliable, we sample Z_t solely to determine whether this shipper is ever allocated a slot; this notion will be formalized shortly in the sequel.

Let $\mathcal{A}_Z = \{t \in \mathcal{A}_{\text{KNAP}} : Z_t = 1\}$ denote the random set of shippers packed by π_{KNAP} for which $Z_t = 1$ and let $\mathcal{A}_f = \mathcal{A}_Z \cup T_{\text{unrel}}$ the final random set of all accepted shippers.

C.1. Analyzing the Competitive Ratio

In what follows, we prove an explicit expression for a lower bound on the expected profit of our proposed policy, which we then directly exploit within a competitive analysis.

A lower bound on the expected profit. To lower bound the expected profit of $\mathcal{A}_f$, we consider a suboptimal slot-allocation mechanism under which only shippers $t \in \mathcal{A}_Z$ can ever be assigned a slot. Under this mechanism, we have

$$\begin{aligned}
\mathbb{E}\left[R(\mathcal{A}_f)\right] \geq \sum_{t\in\mathcal{A}_{\text{KNAP}}} \Pr\left[t \in \mathcal{A}_Z\right] \cdot \Bigg(& \Pr\left[\sum_{\tau\in\mathcal{A}_{\text{KNAP}}\setminus\{t\}} Y_\tau \cdot Z_\tau < m\right] \cdot (p_t + (1-p_t)\cdot d_t) \\
& + \left(1 - \Pr\left[\sum_{\tau\in\mathcal{A}_{\text{KNAP}}\setminus\{t\}} Y_\tau \cdot Z_\tau < m\right]\right) \cdot (1-2p_t)\cdot d_t \Bigg) \\
& + \sum_{t\in\mathcal{A}_{\text{KNAP}}\cap T_{\text{unrel}}} \Pr\left[t \notin \mathcal{A}_Z\right] \cdot (1-2p_t)\cdot d_t + \sum_{t\in T_{\text{unrel}}\setminus\mathcal{A}_{\text{KNAP}}} (1-2p_t)\cdot d_t \\
= (1-\delta) \cdot & \sum_{t\in\mathcal{A}_{\text{KNAP}}} \Pr\left[\sum_{\tau\in\mathcal{A}_{\text{KNAP}}\setminus\{t\}} Y_\tau \cdot Z_\tau < m\right] \cdot (p_t \cdot (1+d_t)) \\
& + \sum_{t\in\mathcal{A}_{\text{KNAP}}} (1-\delta)\cdot(1-2p_t)\cdot d_t + \sum_{t\in\mathcal{A}_{\text{KNAP}}\cap T_{\text{unrel}}} \delta\cdot(1-2p_t)\cdot d_t + \sum_{t\in T_{\text{unrel}}\setminus\mathcal{A}_{\text{KNAP}}} (1-2p_t)\cdot d_t \\
\geq (1-\delta) \cdot & \sum_{t\in\mathcal{A}_{\text{KNAP}}} \Pr\left[\sum_{\tau\in\mathcal{A}_{\text{KNAP}}} Y_\tau \cdot Z_\tau < m\right] \cdot (p_t \cdot (1+d_t)) \\
& + \sum_{t\in\mathcal{A}_{\text{KNAP}}\cap T_{\text{rel}}} (1-\delta)\cdot(1-2p_t)\cdot d_t + \sum_{t\in T_{\text{unrel}}} (1-2p_t)\cdot d_t,
\end{aligned}$$

where the first inequality follows because we evaluate the expected profit under a sub-optimal slot allocation mechanism, further pessimistically assuming that a shipper $t \in \mathcal{A}_Z$ receives a slot only if fewer than m shippers show up among those in $\mathcal{A}_Z \setminus \{t\}$. The second inequality follows from the observation that

$$\Pr\left[\sum_{\tau\in\mathcal{A}_{\mathrm{KNAP}}\setminus\{t\}} Y_\tau\cdot Z_\tau < m\right] \geq \Pr\left[\sum_{\tau\in\mathcal{A}_{\mathrm{KNAP}}} Y_\tau\cdot Z_\tau < m\right],$$

which holds trivially for all $t\in\mathcal{A}_{\mathrm{KNAP}}$.

Establishing the competitive ratio. Our proceeding analysis will make use of the following intermediate lemma, which relates $\mathrm{OPT}_{\mathrm{fluid}}$ to $\mathcal{V}(\mathcal{A}^*)$, the latter of which denotes the optimal value of the optimal offline packing. The proof of this lemma appears in Appendix C.2.

LEMMA 5. $\mathcal{V}(\mathcal{A}^*)+\sum_{t\in T_{\mathrm{unrel}}}(1-2p_t)\cdot d_t \geq (1-\frac{1}{m-1})\cdot \mathrm{OPT}_{\mathrm{fluid}}$

From here, observe that

$$\begin{aligned}
&\left(\mathcal{V}(\mathcal{A}_{\mathrm{KNAP}})+\sum_{t\in T_{\mathrm{unrel}}}(1-2p_t)\cdot d_t\right)-\mathbb{E}\left[R(\mathcal{A}_f)\right]\\
&\qquad\leq\left(1-(1-\delta)\cdot\Pr\left[\sum_{\tau\in\mathcal{A}_{\mathrm{KNAP}}}Y_\tau\cdot Z_\tau<m\right]\right)\cdot\sum_{t\in\mathcal{A}_{\mathrm{KNAP}}}p_t\cdot(1+d_t)\\
&\qquad\qquad+\sum_{t\in\mathcal{A}_{\mathrm{KNAP}}\cap T_{\mathrm{rel}}}\delta\cdot\underbrace{(1-2p_t)\cdot d_t}_{<=0}\\
&\qquad\leq\left(1-(1-\delta)\cdot\Pr\left[\sum_{\tau\in\mathcal{A}_{\mathrm{KNAP}}}Y_\tau\cdot Z_\tau<m\right]\right)\cdot\sum_{t\in\mathcal{A}_{\mathrm{KNAP}}}p_t\cdot(1+d_t)\\
&\qquad\leq\left(1-(1-\delta)\cdot\Pr\left[\sum_{\tau\in\mathcal{A}_{\mathrm{KNAP}}}Y_\tau\cdot Z_\tau<m\right]\right)\cdot 2\mathcal{V}(\mathcal{A}_{\mathrm{KNAP}}),
\end{aligned}\tag{7}$$

where the final inequality follows since

$$\begin{aligned}
2\mathcal{V}(\mathcal{A}_{\mathrm{KNAP}}) &= 2\cdot\left(\sum_{t\in\mathcal{A}_{\mathrm{KNAP}}}p_t\cdot(1+d_t)+\sum_{t\in\mathcal{A}_{\mathrm{KNAP}}\cap T_{\mathrm{rel}}}(1-2p_t)\cdot d_t\right)\\
&=\sum_{t\in\mathcal{A}_{\mathrm{KNAP}}}p_t\cdot(1+d_t)+\sum_{t\in\mathcal{A}_{\mathrm{KNAP}}\cap T_{\mathrm{unrel}}}p_t\cdot(1+d_t)+\sum_{t\in\mathcal{A}_{\mathrm{KNAP}}\cap T_{\mathrm{rel}}}\underbrace{p_t+2d_t-3p_td_t}_{\geq 0}\\
&\geq\sum_{t\in\mathcal{A}_{\mathrm{KNAP}}}p_t\cdot(1+d_t).
\end{aligned}$$

Together, this implies that

$$\begin{aligned}
&\mathbb{E}\left[R(\mathcal{A}_f)\right]\\
&\quad\geq\left(1-2\left(1-(1-\delta)\cdot\Pr\left[\sum_{\tau\in\mathcal{A}_{\mathrm{KNAP}}}Y_\tau\cdot Z_\tau<m\right]\right)\right)\cdot\left(\mathcal{V}(\mathcal{A}_{\mathrm{KNAP}})+\sum_{t\in T_{\mathrm{unrel}}}(1-2p_t)\cdot d_t\right)\\
&\quad\geq\left(1-2\left(1-(1-\delta)\cdot\Pr\left[\sum_{\tau\in\mathcal{A}_{\mathrm{KNAP}}}Y_\tau\cdot Z_\tau<m\right]\right)\right)\cdot 0.59\left(1-\frac{1}{m}\right)\cdot\left(\mathcal{V}(\mathcal{A}^*)+\sum_{t\in T_{\mathrm{unrel}}}(1-2p_t)\cdot d_t\right)
\end{aligned}$$

$$\geq \left(1 - 2\left(1 - (1-\delta)\cdot \Pr\left[\sum_{\tau\in\mathcal{A}_{\text{KNAP}}} Y_\tau \cdot Z_\tau < m\right]\right)\right)\cdot 0.59\left(1 - \frac{1}{m-1}\right)^2 \text{OPT}_{\text{fluid}}$$

$$= \left(2(1-\delta)\cdot \Pr\left[\sum_{\tau\in\mathcal{A}_{\text{KNAP}}} Y_\tau \cdot Z_\tau < m\right] - 1\right)\cdot 0.59\left(1 - \frac{1}{m-1}\right)^2 \text{OPT}_{\text{fluid}} \tag{8}$$

where the first inequality follows from (7), and the second follows from Theorem 4, noting that within our reduction, we have that $L = 1$, $U = 2$ and $\epsilon = \frac{1}{m}$. The last inequality follows from Lemma 5. From here, we note that since Y_τ and Z_τ are clearly independent Bernoulli random variables, it is the case that $Y_\tau \cdot Z_\tau \sim \text{Bernoulli}((1-\delta)\cdot p_t)$. As a result we have that

$$\Pr\left[\sum_{\tau\in\mathcal{A}_{\text{KNAP}}} Y_\tau \cdot Z_\tau < m\right] \geq \Pr\left[\sum_{\tau\in\mathcal{A}_{\text{KNAP}}} Y_\tau \cdot Z_\tau < (1-\delta)(1+\delta)\cdot m\right]$$
$$\geq 1 - \exp(-\frac{m\cdot(1-\delta)\delta^2}{2+\delta}),$$

where the second inequality follows from a standard application of the multiplicative Chernoff bound, along with the fact that $\mathbb{E}[\sum_{\tau\in\mathcal{A}_{\text{KNAP}}} Y_\tau \cdot Z_\tau] = (1-\delta)\cdot m$. From here, plugging in $\delta = \sqrt{\frac{\log m}{m}}$, we get that

$$(8) \geq \left(2\left(1 - \sqrt{\frac{\log m}{m}}\right)\cdot\left(1 - m^{-\frac{1}{2}+o(1)}\right) - 1\right)\cdot 0.59\left(1 - \frac{1}{m-1}\right)^2 \text{OPT}_{\text{fluid}},$$

as desired.

C.2. Proof of Lemma 5

We begin by presenting an intermediate claim, which reveals a handful of useful properties regarding the optimal solution to (Fluid IP). The proof of this claim appears at the end of this section.

CLAIM 8. *There exists an optimal solution (x^*, y^*, z^*) to (Fluid IP) that satisfies the following three properties:*

(i) There is at most one shipper $t \in [T]$ such that $p_t > y_t^* > 0$.
(ii) If there exists a shipper $t \in [T]$ such that $p_t > y_t^* > 0$, then $\text{OPT}_{\text{Fluid}} \geq m - 1$.
(iii) If $y_t^* = 0$, then $x_t^* = 1$ if and only if $t \in T_{\text{unrel}}$.

Proof of lemma. Let $\mathcal{Y}^* = t \in [T] : y_t^* > 0$ and $\bar{\mathcal{Y}}^* = t \in [T] : y_t^* = p_t$. In addition, letting $\mathcal{A}^* = t \in [T] : x_t^* = 1$ allows us to express the optimal objective value of (Fluid IP) as

$$\text{OPT}_{\text{Fluid}} = \sum_{t\in\bar{\mathcal{Y}}^*} p_t + (1-p_t)\cdot d_t + \sum_{t\in\mathcal{Y}^*\setminus\bar{\mathcal{Y}}^*} y_t\cdot(1-d_t) + (1-2p_t)\cdot d_t + \sum_{t\in\mathcal{A}^*\setminus\mathcal{Y}^*} (1-2p_t)\cdot d_t$$

$$= \sum_{t\in\bar{\mathcal{Y}}^*\cap T_{\mathrm{rel}}} p_t + (1-p_t)\cdot d_t + \sum_{t\in\bar{\mathcal{Y}}^*\cap T_{\mathrm{unrel}}} p_t\cdot(1+d_t) + \sum_{t\in T_{\mathrm{unrel}}}(1-2p_t)\cdot d_t$$
$$\sum_{t\in\mathcal{Y}^*\setminus\bar{\mathcal{Y}}^*} y_t\cdot(1-d_t) + \mathbf{1}_{t\in T_{\mathrm{rel}}}\cdot(1-2p_t)\cdot d_t.$$

The second inequality above uses property (iii) of Claim 8, which gives that $\mathcal{A}^*\setminus\mathcal{Y}^*\subseteq T_{\mathrm{unrel}}$. Now, consider the solution $\bar{\mathcal{Y}}^*$ within the offline knapsack problem, i.e., we select all shippers $t\in\bar{\mathcal{Y}}^*$. This solution is feasible by constraint (1) in (Fluid IP), and its total value satisfies

$$\begin{aligned}\mathcal{V}(\bar{\mathcal{Y}}^*) + \sum_{t\in T_{\mathrm{unrel}}}(1-2p_t)\cdot d_t &= \sum_{t\in\bar{\mathcal{Y}}^*\cap T_{\mathrm{rel}}} p_t + (1-p_t)\cdot d_t + \sum_{t\in\bar{\mathcal{Y}}^*\cap T_{\mathrm{unrel}}} p_t\cdot(1+d_t) + \sum_{t\in T_{\mathrm{unrel}}}(1-2p_t)\cdot d_t\\ &\geq \mathrm{OPT}_{\mathrm{Fluid}} - \sum_{t\in\mathcal{Y}^*\setminus\bar{\mathcal{Y}}^*} y_t\cdot(1-d_t)\\ &\geq (1-\frac{1}{m-1})\cdot\mathrm{OPT}_{\mathrm{Fluid}},\end{aligned}$$

as desired. The last inequality follows because if $\mathcal{Y}^*\setminus\bar{\mathcal{Y}}^*\neq\emptyset$, then by property (i) of Claim 8, we have $|\mathcal{Y}^*\setminus\bar{\mathcal{Y}}^*| = 1$ and so

$$\sum_{t\in\mathcal{Y}^*\setminus\bar{\mathcal{Y}}^*} y_t\cdot(1-d_t)\leq 1\leq\frac{1}{m-1}\cdot\mathrm{OPT}_{\mathrm{Fluid}}$$

where the final inequality follows by property (ii) of Claim 8.

Proof of Claim 8 - property (i). For any optimal solution (x^*,y^*,z^*), define $\mathcal{Y}^* = \{t\in[T]: y_t^* > 0\}$ and $\bar{\mathcal{Y}}^* = \{t\in[T]: y_t^* = p_t\}$. If $|\mathcal{Y}^*\setminus\bar{\mathcal{Y}}^*| > 1$, we construct an alternative optimal solution for which $|\mathcal{Y}^*\setminus\bar{\mathcal{Y}}^*|$ is reduced by at least one. This procedure can then be applied iteratively until an optimal solution satisfying property (i) is obtained.

Specifically, consider two indices $t,t'\in\mathcal{Y}^*\setminus\bar{\mathcal{Y}}^*$ and assume without loss of generality that $d_t\geq d_{t'}$. Let $\epsilon=\min\{y_{t'}^*, p_t - y_t^*\}$ and define an alternative solution $(\hat{x},\hat{y},\hat{z})$ that coincides with (x^*,y^*,z^*) except that $\hat{y}_t = y_t^*+\epsilon$, $\hat{z}_t = z_t^*-\epsilon$, and $\hat{y}_{t'} = y_{t'}^*-\epsilon$, $\hat{z}_{t'} = z_{t'}^*+\epsilon$. This solution is clearly feasible and results in an objective change of $\epsilon\cdot(d_t-d_{t'})\geq 0$. Moreover, by construction, we have

$$|\{t\in[T]:\hat{y}_t>0\}\setminus\{t\in[T]:\hat{y}_t=p_t\}| < |\mathcal{Y}^*\setminus\bar{\mathcal{Y}}^*|$$

by construction and as desired.

Proof of Claim 8 - property (ii). Assume by way of contradiction that there exists shipper τ such that $p_\tau > y_\tau^* > 0$ and $\mathrm{OPT}_{\mathrm{Fluid}} < m-1$. To establish the contradiction, observe that if $\mathrm{OPT}_{\mathrm{Fluid}} < m-1$, then it must be the case that $\sum_{t\in[T]} y_t^* < m-1$. Consequently, increasing y_τ^* by any $0<\epsilon\leq p_\tau - y_\tau^*$ and decreasing z_τ^* by ϵ preserves feasibility and strictly improves the objective value, thereby contradicting the optimality of the proposed solution.

Proof of Claim 8 - property (iii). This property follows trivially from the fact that $(1-2p_t)\cdot d_t \geq 0$, only for reliable shippers.

Appendix D: Illustrative Examples

D.1. An all-knowing clairvoyant

The following example shows that if the clairvoyant had access to the show-up decision of each shipper, then it could garner an expected profit that is arbitrarily larger than $\mathrm{OPT}_{\mathrm{fluid}}$.

EXAMPLE 3. Consider a setting with $m = 1$, where the deposit sequence is $(1,1)$ followed by infinite copies of the deposit tuple $(1, 0.9)$. A clairvoyant that knows which shippers show up will garner infinite profit in expectation; from each shipper beyond the first, the benchmark earns an expected profit of at least $(1-0.9)*1 = 0.1$ by simply accepting each shipper that is a no-show. It is straightforward to see that for this sequence, we have $\mathrm{OPT}_{\mathrm{fluid}} = 1$

D.2. Additional details for Example 2

We first describe how we determine $q(p)$ for fixed $p \in [0.5, 1]$. From here, we formulate the optimization problem to determine the worst-case p and the resulting upper bound on the best achievable competitive ratio of any deterministic policy. In what follows, we use $R(\mathcal{A};m)$ to denote the random profit of acceptance set $\mathcal{A}$ given that there are m available slots.

Determining $\boldsymbol{q(p)}$. In relation to Stream 2, let $\mathcal{A}(q)$ denote an acceptance set that includes q copies of the deposit tuple $(1-\epsilon, p)$. We define $q(p) = \max\{q \in \mathbb{Z}_+ : \mathbb{E}[R(\mathcal{A}(q);m) - R(\mathcal{A}(q-1);m)] > 0\}$ to be the maximum number of shippers such that if the adversary sent only Stream 2, it would be optimal to accept the entire stream. It will always be the case that $\mathbb{E}[R(\mathcal{A}(q(p));m)] \geq m$.

Computing the CR upper bound. Any deterministic policy is fully specified by the number of shippers $\alpha_1, \alpha_2 \in \mathbb{Z}_+$ accepted from the first two streams. Note that since the shippers from Stream 3 are unreliable, it is always optimal to accept all of them. For any such policy, we derive its competitive ratio based on the deposit sequence selected by the adversary among the three under consideration:

- Case 1 - Just Stream 1: Given α_1, the adversary can decide to stop the arrival process. If it chooses to do so, the proposed policy earns a competitive ratio of

$$\mathrm{CR}_1(\alpha_1) = \frac{\alpha_1}{m},$$

  since the clairvoyant will accept the entire Stream 1 and earn a profit of m.

- Case 2 - Streams 1 and 2: In this case, the α_1 accepted shippers from Stream 1 consume this many slots (since they show up with certainty), leaving $m-\alpha_1$ slots left for the shippers from Stream 2. Accordingly, from Stream 2, the policy earns $\mathbb{E}\left[R(\mathcal{A}(\alpha_2);m-\alpha_1)\right]$, leading to a competitive ratio of

$$\mathrm{CR}_2\left(\alpha_1,\alpha_2,p\right)=\frac{\alpha_1+\mathbb{E}\left[R(\mathcal{A}(\alpha_2);m-\alpha_1)\right]}{\mathbb{E}\left[R(\mathcal{A}(q(p));m)\right]},$$

as the clairvoyant with reject all shippers from Stream 1 and accept everyone from Stream 2.
- Case 3 - Streams 1,2 and 3: In this third case, the clairvoyant will reject all of Streams 1 and 2 and accept all of Stream 3 to earn a profit of $2m$. All-in-all, noting that we can represent Stream 3 as $\{t\in[T]:t>m+q(p)\}$, the third competitive ratio is

$$\mathrm{CR}_3\left(\alpha_1,\alpha_2,p\right)=\frac{\alpha_1+\mathbb{E}\left[R(\mathcal{A}(\alpha_2)\cup\{t\in[T]:t>m+q(p)\};m-\alpha_1)\right]}{2m},$$

The best possible competitive ratio plotted in Figure 1 is the objective of the following optimization problem

$$\min_{p\in[0.5,1]}\max_{\alpha_1,\alpha_2}\min\left\{\mathrm{CR}_1\left(\alpha_1\right),\mathrm{CR}_2\left(\alpha_1,\alpha_2,p\right),\mathrm{CR}_3\left(\alpha_1,\alpha_2,p\right)\right\},$$

which we solve by discretizing p and enumerating over all possible α_1 and α_2.

Appendix E: Additional Numerical Experiments

E.1. Results - model misspecification

In this section, we evaluate the robustness of our proposed policies under model misspecification. Specifically, we test the RT policy under the coupled shipper arrival model, and both KP and the GUM policy under the independent shipper arrival model when the liner incorrectly specifies the functional relationship between deposit and show-up probability. More specifically, we consider all combinations of model misspecification, where each case is distinguished by a presumed model $\mathcal{P}_? \in \{\text{Linear,Concave,Convex}\}$, and a ground-truth model $\mathcal{P}\in\{\text{Linear,Concave,Convex}\}$. When $\mathcal{P}_?\neq\mathcal{P}$, the model is misspecified. In this section, we only measure the implications of model misspecification on overall performance as measured by the optimality gaps.

The RT policy under coupled arrivals. Table 3 shows the performance of the RT policy under the coupled arrival setting. The latter three column blocks report the assumed show-up probability model $\mathcal{P}_?$; for each (potentially misspecified) model, we report the optimality gap of the RT policy as well as the average number of accepted shippers per slot.

Parameters			Linear		Concave		Convex	
$(T, M, \mathcal{P})$			Opt. Gap (%)	# Accepted	CR (%)	# Accepted	CR (%)	# Accepted
100	10	Concave	92.56	1.42	95.21	1.17	84.62	1.64
100	10	Convex	90.58	1.43	84.39	1.20	93.85	1.67
100	10	Linear	92.79	1.44	89.77	1.19	89.79	1.67
100	25	Concave	92.99	1.42	95.38	1.18	87.02	1.60
100	25	Convex	92.40	1.42	85.88	1.18	95.25	1.60
100	25	Linear	93.59	1.43	90.23	1.18	91.79	1.61
100	50	Concave	94.87	1.38	96.27	1.17	93.29	1.44
100	50	Convex	96.73	1.38	90.82	1.18	98.61	1.45
100	50	Linear	96.23	1.38	92.79	1.17	96.39	1.44
250	10	Concave	92.23	1.44	95.03	1.19	83.72	1.67
250	10	Convex	90.60	1.44	84.10	1.19	93.69	1.67
250	10	Linear	92.86	1.45	89.63	1.19	89.79	1.67
250	25	Concave	92.37	1.44	95.10	1.19	84.40	1.66
250	25	Convex	90.84	1.45	84.10	1.19	93.87	1.67
250	25	Linear	92.86	1.44	89.76	1.18	89.96	1.65
250	50	Concave	92.47	1.44	95.17	1.19	85.60	1.63
250	50	Convex	91.65	1.44	85.01	1.19	94.64	1.64
250	50	Linear	93.23	1.43	90.00	1.19	90.85	1.63
500	10	Concave	92.40	1.43	95.03	1.18	83.65	1.67
500	10	Convex	90.61	1.45	83.79	1.19	93.60	1.67
500	10	Linear	92.67	1.43	89.55	1.17	89.43	1.66
500	25	Concave	92.36	1.45	95.09	1.19	84.09	1.66
500	25	Convex	90.53	1.45	83.92	1.19	93.54	1.68
500	25	Linear	92.64	1.45	89.45	1.19	89.41	1.68
500	50	Concave	92.28	1.44	95.08	1.19	84.05	1.67
500	50	Convex	90.71	1.44	84.02	1.19	93.80	1.67
500	50	Linear	92.74	1.44	89.69	1.18	89.70	1.66

Table 3 Performance of RT policy with misspecification

Figure 3 presents a heatmap of the average optimality gap under model misspecification. The diagonal entries correspond to correctly specified models, while the off-diagonal entries represent performance under misspecification. Overall, the RT policy demonstrates strong robustness to model misspecification, with optimality gaps remaining high across all scenarios.

When the true show-up probability model is Linear or Convex, misspecification leads to only moderate performance degradation. Across these cases, assuming an incorrect model results in optimality gap losses typically below 5%, and often substantially smaller, particularly for larger values of m. This indicates that the RT policy is relatively insensitive to moderate errors in the assumed show-up probability structure.

A notably different pattern emerges when the true model is Concave. In this case, correct specification yields the highest competitive ratios that are consistently above 95%, while misspecifying the model as Linear or Convex leads to a more significant drop in performance. This decline is accompanied by a substantial increase in the average number of accepted shippers per slot. As shown in Table 3, when the true model is Concave, assuming a Linear or Convex model results in approximately 0.20–0.45 additional acceptances per slot relative to the correctly specified Concave case. This behavior can be attributed to the interaction between model curvature and the RT policy's threshold structure. Concave show-up functions generate higher show-up probabilities and larger effective rewards, which naturally encourage a more conservative booking strategy. When the model is misspecified as Linear or Convex, the policy underestimates these effects and accepts more shippers, increasing congestion and reducing overall performance. Since the RT thresholds are computed under a worst-case reward-penalty formulation, this over-acceptance effect becomes particularly pronounced when the true environment is Concave, revealing the cost of excessive conservatism under correct specification.

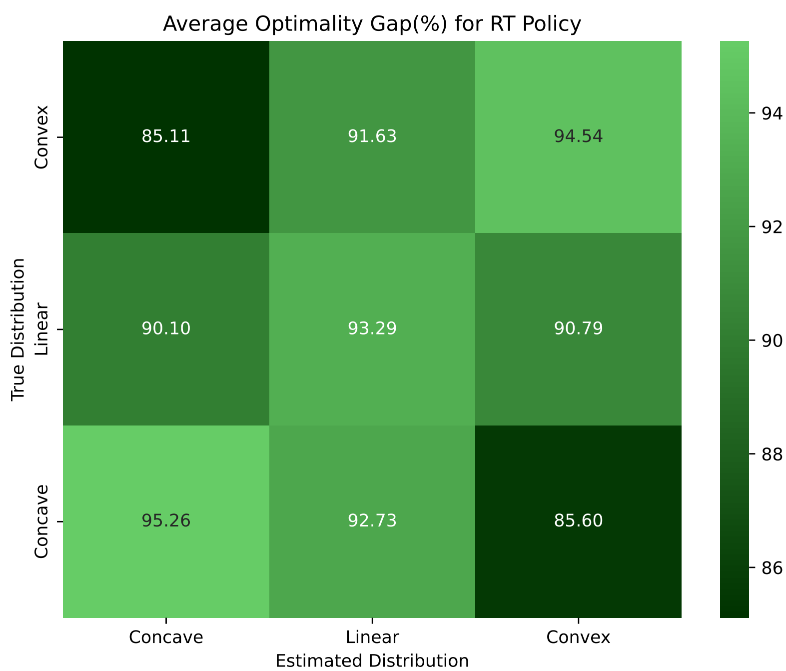

Figure 3 Model Misspecification for Coupled Case

Independent arrival setting. Table 4 shows the performance of the GUM and KP policies in the independent arrival setting. The fourth to sixth columns report the optimlaity gap achieved by the GUM policy when the decision-maker assumes that the relationship between deposit and show-up probability is Linear, Concave, or Convex, respectively. Similarly, the seventh to ninth columns present the optimality gap of the KP policy under the same model misspecification scenarios. Figure 4 visualizes the data in Table 4 by presenting heatmaps of the average optimality gap for both KP and the GUM policy under model misspecification, consistent in style with Figure 3.

Parameters			GUM			KP		
$(T, M, \mathcal{P})$			Linear	Concave	Convex	Linear	Concave	Convex
100	10	Concave	96.47	93.21	93.37	83.35	65.38	96.05
100	10	Convex	72.10	62.29	81.80	53.59	42.58	72.64
100	10	Linear	88.14	78.85	92.54	68.22	54.14	88.57
100	25	Concave	97.36	95.35	93.63	80.67	62.99	97.37
100	25	Convex	75.12	64.48	86.06	53.00	41.45	73.08
100	25	Linear	91.26	80.59	94.55	66.54	52.06	90.24
100	50	Concave	97.92	96.68	93.99	80.25	62.32	98.23
100	50	Convex	80.38	68.92	93.12	55.81	43.54	77.27
100	50	Linear	94.31	82.61	96.80	67.28	52.32	92.46
250	10	Concave	96.43	93.34	93.30	82.96	65.45	95.85
250	10	Convex	71.91	62.24	81.51	53.38	42.20	72.39
250	10	Linear	88.18	78.81	92.58	68.38	53.86	88.61
250	25	Concave	97.21	95.30	93.33	80.64	62.80	97.27
250	25	Convex	73.68	63.18	84.50	51.97	40.62	71.77
250	25	Linear	90.79	80.23	94.05	66.12	51.75	89.69
250	50	Concave	97.37	96.29	93.37	79.88	61.99	97.68
250	50	Convex	74.98	64.20	86.72	52.03	40.64	72.08
250	50	Linear	92.24	80.78	94.76	65.72	51.15	90.47
500	10	Concave	96.47	93.33	93.16	83.40	65.42	95.96
500	10	Convex	71.80	61.99	81.22	53.55	41.97	72.18
500	10	Linear	88.06	78.66	92.58	68.38	53.71	88.66
500	25	Concave	97.24	95.24	93.47	80.50	62.89	97.29
500	25	Convex	73.33	63.00	84.28	51.82	40.44	71.56
500	25	Linear	90.63	80.06	94.06	66.24	51.72	89.66
500	50	Concave	97.37	96.17	93.52	79.71	61.95	97.68
500	50	Convex	74.01	63.47	85.76	51.42	40.11	71.35
500	50	Linear	91.95	80.58	94.37	65.60	51.03	90.09

Table 4 Performance of GUM and KP policy with misspecification

Both policies continue to exhibit stable and strong optimality gaps under model misspecification. In the vast majority of configurations, the GUM policy outperforms the KP policy, often by a substantial margin. This dominance holds both under correct specification and across most misspecification scenarios, highlighting the robustness of the distribution-agnostic design of GUM.There is, however, a notable exception to this pattern. When the true show-up probability model is Concave but the decision-maker incorrectly assumes a Convex model, the KP policy can outperform GUM. As shown in Table 4, this is one of the few settings in which KP attains a higher optimality gap than GUM. We conjecture that this reversal is driven by the interaction between misspecification and the acceptance behavior of KP. Under correct specification, KP tends to be overly conservative in the acceptance decision, leading to under-acceptance of shippers and suboptimal utilization. When the model is misspecified as Convex, this conservatism is partially offset, resulting in more aggressive acceptance and improved performance.

Fixing the estimated model $\mathcal{P}_? \in \{\text{Concave}, \text{Convex}, \text{Linear}\}$, we further observe a systematic dependence of performance on the true underlying model. For both policies—but most prominently for GUM—the optimality gap is consistently lowest when the true model is Convex and highest when the true model is Concave. This pattern is robust across all values of T and m. The Convex model induces lower show-up probabilities relative to the Linear and Concave models, which leads to conservative acceptance decisions and under-utilization of available capacity. Consequently, under any fixed estimated model, when the true model is Convex, the total realized show-up probability of accepted shippers frequently falls below the capacity level m, resulting in inferior competitive performance.

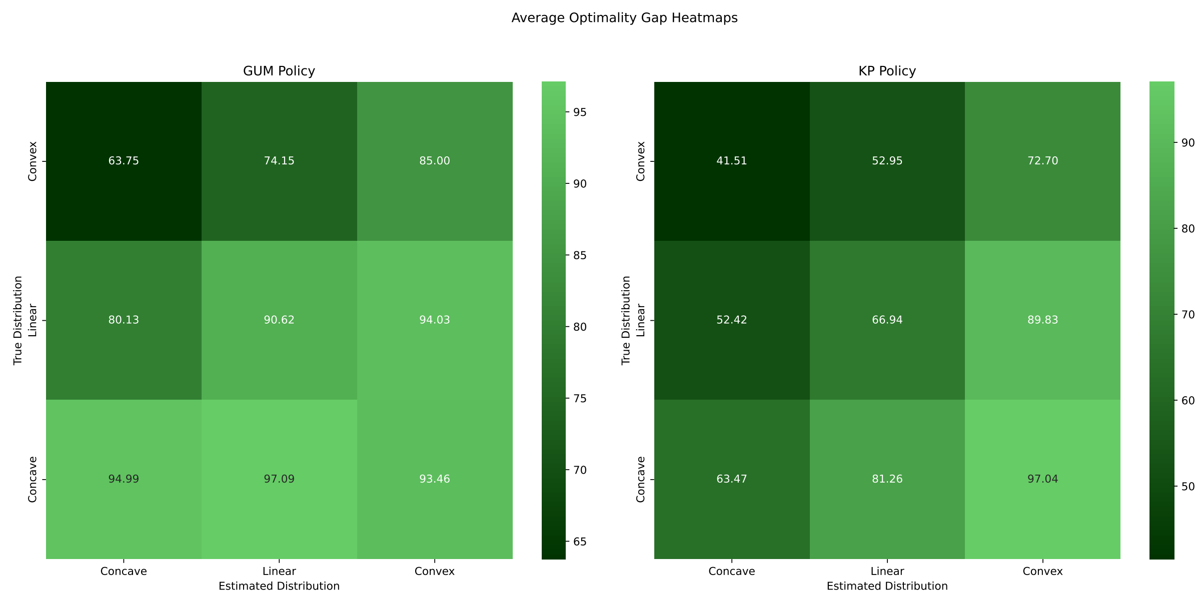

Figure 4 Model Misspecification for Independent Case

E.2. Results - hybrid demand model

In this section, we evaluate the performance of the RT, KP, and GUM policies under a hybrid demand model. In this model, the show-up behavior of a subset of shippers is correlated through the attractiveness of a common outside option (e.g., a spot market), while the show-up behavior of the remaining shippers is driven by their own individual outside option and is therefore independent of the decisions of other accepted shippers.

The hybrid model. In essence, our hybrid model is designed to capture a mixture of independent and coupled show-up behavior. To this end, let $S \sim \text{Unif}(0,1)$ denote a common market factor, and let $\{U_t\}_{t=1}^{T}$ be i.i.d. with $U_t \sim \text{Unif}(0,1)$, independent of S. For each arrival t, we draw an indicator $I_t \sim \text{Bern}(\lambda)$ independently across t and independent of S and U_t, which determines the effective “type” of shipper t, i.e., if they will be an independent or coupled show-up shipper. Specifically, the random attractiveness of the outside option for shipper t is denoted by

$$X_t = \begin{cases} S, & \text{if } I_t = 1, \\ U_t, & \text{if } I_t = 0, \end{cases} \qquad t = 1, \ldots, T.$$

A shipper arriving at time t with deposit $d_t \in [0,1]$ shows up if

$$1 \le d_t + X_t.$$

The resulting show-up probability is

$$p(d) = \Pr(X_t \ge 1 - d) = 1 - F_X(1-d),$$

where F_X is the CDF of X_t. Thus, with probability λ the outside option is driven by the common factor S, thereby coupling the show-up decisions of all shippers with $I_t = 1$, while with probability $1-\lambda$ a shipper exhibits independent show-up behavior. In particular, $\lambda = 0$ corresponds to the independent show-up setting, whereas $\lambda = 1$ recovers the fully coupled show-up setting.

Computational set-up. We set the arrival horizon to $T = 300$ and consider capacities $m \in \{50, 100, 200\}$ and mixture levels $\lambda \in \{0.25, 0.5, 0.75\}$. For each sample path, deposits are drawn independently as $d_t \sim \text{Unif}(1/2, 1)$, ensuring all arrivals are reliable. We study two arrival-order settings: `random`, in which deposits arrive in random order, and `increase`, in which deposits are sorted increasingly so that higher deposits tend to appear later in the horizon. For each parameter configuration and each benchmark policy $\Pi \in \{\text{GUM}, \text{KP}, \text{RT}\}$, we generate 500 deposit sequences and estimate the expected profit of each policy's acceptance set using 1000 Monte Carlo trials. The optimal gap of each policy π is estimated as

$$\text{GAP}(\Pi) = \frac{\mathbb{E}[\text{Rev}(\Pi)]}{\text{OPT}_{\text{fluid}}},$$

noting that $\text{OPT}_{\text{fluid}}$ continues to serve as a valid upper bound on the achievable expected revenue under all arrival dependence regimes considered—independent, fully dependent, and mixed—since its derivation does not require any independence assumptions on the arrival process; see the proof of Claim 2.

Results. Table 5 reports the performance of the RT, GUM, and KP policies across different mixture levels λ, capacity values m, and arrival-order settings. The first three columns describe the experimental parameters, while the middle block reports the average percent optimality gap achieved by each policy. The last two columns present the percentage improvement difference between GUM and RT and between KP and RT, respectively, where positive values indicate that the corresponding benchmark outperforms RT.

Parameters			GAP(Π) (%)			Diff with RT	
Mode	λ	m	GUM	KP	RT	GUM	KP
increase	0.25	50	89.13	78.99	82.88	6.25	-3.89
increase	0.25	100	90.85	76.37	87.46	3.39	-11.09
increase	0.25	200	93.48	71.84	89.03	4.45	-17.19
increase	0.50	50	78.66	74.15	78.37	0.29	-4.22
increase	0.50	100	81.88	75.59	80.66	1.22	-5.07
increase	0.50	200	86.99	71.83	85.50	1.49	-13.67
increase	0.75	50	67.97	65.67	71.99	-4.02	-6.32
increase	0.75	100	72.81	69.89	73.34	-0.53	-3.45
increase	0.75	200	80.45	71.81	79.76	0.69	-7.95
random	0.25	50	80.26	73.62	75.19	5.07	-1.57
random	0.25	100	84.45	76.86	78.21	6.24	-1.35
random	0.25	200	92.16	80.09	81.19	10.97	-1.10
random	0.50	50	75.58	72.48	73.57	2.01	-1.09
random	0.50	100	79.21	76.14	76.86	2.35	-0.72
random	0.50	200	86.40	80.18	81.50	4.90	-1.32
random	0.75	50	70.30	68.83	69.96	0.34	-1.13
random	0.75	100	73.81	72.11	73.35	0.46	-1.24
random	0.75	200	80.42	78.16	79.10	1.32	-0.94

Table 5 performance for misspecified arrival dependency.

Several systematic patterns emerge. First, we observe that one or both of RT and GUM outperform KP across all parameter combinations, consistent with our findings in experiments where show-up behavior was either purely independent or fully coupled. Second, as expected, when λ is small and most shippers exhibit independent show-up behavior, the GUM policy performs best. In contrast, when $\lambda = 0.75$ and show-up behavior is largely coupled, the RT policy matches or exceeds the performance of GUM, particularly for smaller values of m.

Second, arrival order remains a decisive factor in relative policy performance. Under the `increase` mode, RT consistently outperforms KP across all values of λ and m, while its performance relative to GUM depends on both capacity and the strength of dependence. In particular, GUM retains an advantage for small to moderate values of m when λ is low, but this advantage diminishes as either m or λ increases. By contrast, under `random` ordering, RT generally underperforms GUM and only marginally outperforms KP, indicating that the effectiveness of RT as a threshold-based policy is closely tied to structured arrival sequences in which higher-value shippers tend to arrive later.

Finally, we observe that the performance of the KP policy may deteriorate as m increases from 100 to 200 under the `increase` arrival sequence. We conjecture that this occurs because, although KP raises its acceptance threshold after each acceptance, it continues to admit early arrivals in an increasing sequence, leading to systematic underutilization of later high-value arrivals. Moreover, its overly conservative behavior results in acceptance levels below the fluid benchmark, causing KP to forgo opportunities to accept the later arrival of higher-deposit shippers.

References for Appendix

Alaei S, Hajiaghayi M, Liaghat V (2012) Online prophet-inequality matching with applications to ad allocation. *Proceedings of the 13th ACM Conference on Electronic Commerce*, 18–35.

Chakrabarty D, Zhou Y, Lukose R (2008) Online knapsack problems. *Workshop on internet and network economics (WINE)*, 1–9.

Ma W, Simchi-Levi D, Zhao J (2021) Dynamic pricing (and assortment) under a static calendar. *Management Science* 67(4):2292–2313.