Electronic Companion for Online Fair Allocation of Perishable Resources

Appendix A: Section 3 Omitted Proofs

A.1. Proof of Theorem 2

Proof. We first argue that offset-expiry implies feasibility of B/N. Consider the allocation schedule which allocates goods in increasing order of perishing time (breaking ties arbitrarily), and is such that $X_{t,\theta} = B/N$ for all t, θ, as long as there are resources remaining. Noting that $(B/N)N_{<t}$ is precisely the cumulative allocation at the beginning of round t, this implies that we allocate (weakly) more than the number of goods with perishing time before round t (i.e. $P_{<t}$). Since we allocate goods in increasing order of perishing time, this also implies that no unit ever perishes under this sequence of allocations. Thus, the total allocation by the end of the horizon is $\frac{B}{N} \cdot N = B$, implying that B/N is feasible.

We now argue that offset-expiry is necessary for B/N to be feasible. To see this, consider the first period $t \geq 2$ for which $P_{<t}/B > N_{<t}/N$ (i.e., by the end of period $t-1$, there existed some unallocated goods that had perished). Then, the remaining budget at the start of period t for any algorithm, denoted by B_t^{alg}, is:

$$B_t^{alg} \leq B - P_{<t} < B - N_{<t} \cdot \frac{B}{N} = N_{\geq t} \cdot \frac{B}{N},$$

which implies that the remaining budget does not suffice to allocate B/N to all arrivals from t onwards. Hence, B/N is not feasible. □

Appendix B: Section 4 Omitted Proofs

B.1. Tightness of bounds

Consider the random problem instance which achieves the lower bounds of Theorem 1 with probability 1/2, and the lower bounds of Theorem 3 with probability 1/2. Putting these two bounds together, we have:

$$\mathbb{E}[\Delta_{EF}] \gtrsim \mathcal{L}^{\mathsf{perish}} + 1/\sqrt{T}.$$

By Theorem 4, our algorithm achieves $\mathbb{E}[\Delta_{EF}] \lesssim \max\{L_T, \mathcal{L}^{\mathsf{perish}} + 1/\sqrt{T}\}$. Letting $L_T \lesssim \mathcal{L}^{\mathsf{perish}} + 1/\sqrt{T}$ then, our algorithm achieves this lower bound. We now argue that our algorithm is tight with respect to efficiency in this regime. Suppose $L_T = 0$. By Theorem 1 and Theorem 3, any online algorithm incurs:

$$\mathbb{E}[\Delta_{\mathit{efficiency}}] \gtrsim T\mathcal{L}^{\mathsf{perish}} + \sqrt{T},$$

which is achieved by our algorithm.

Consider now the regime in which $\Delta_{EF} = L_T$, i.e., $L_T \gtrsim \mathcal{L}^{\mathsf{perish}} + 1/\sqrt{T}$. Again, randomizing between the two lower bounds, we have:

$$\mathbb{E}[\Delta_{\mathit{efficiency}}] \gtrsim T\mathcal{L}^{\mathsf{perish}} + \min\{\sqrt{T}, L_T^{-1}\}. \tag{EC.1}$$

Case 1: $L_T^{-1} + \sqrt{T\mathcal{L}^{\mathsf{perish}} L_T^{-1}} \gtrsim \sqrt{T}$. Here, our algorithm achieves $\mathbb{E}[\Delta_{\mathit{efficiency}}] \lesssim \sqrt{T} + T\mathcal{L}^{\mathsf{perish}}$.

If $L_T^{-1} \gtrsim \sqrt{T}$, we achieve the bound in (EC.1). Suppose now that $L_T^{-1} = o(\sqrt{T})$. Then, (EC.1) implies that $\mathbb{E}[\Delta_{\mathit{efficiency}}] \gtrsim T\mathcal{L}^{\mathsf{perish}} + L_T^{-1}$. We argue that, if $L_T^{-1} = o(\sqrt{T})$, then in this case $T\mathcal{L}^{\mathsf{perish}} \gtrsim \sqrt{T}$. $T\mathcal{L}^{\mathsf{perish}}$ then dominates both the lower bound in (EC.1), as well as our upper bound, which gives us tightness.

Case 2: $L_T^{-1} + \sqrt{T\mathcal{L}^{\mathsf{perish}} L_T^{-1}} \lesssim \sqrt{T}$. Here, our algorithm achieves $\mathbb{E}[\Delta_{\mathit{efficiency}}] \lesssim L_T^{-1} + \sqrt{T\mathcal{L}^{\mathsf{perish}} L_T^{-1}} + T\mathcal{L}^{\mathsf{perish}}$. Since $L_T^{-1} \lesssim \sqrt{T}$, (EC.1) reduces to $\mathbb{E}[\Delta_{\mathit{efficiency}}] \gtrsim T\mathcal{L}^{\mathsf{perish}} + L_T^{-1}$. It is easy to check that $\sqrt{T\mathcal{L}^{\mathsf{perish}} L_T^{-1}} \lesssim \max\{L_T^{-1}, T\mathcal{L}^{\mathsf{perish}}\}$, which completes the tightness argument.

B.2. Section 4.2 Omitted Proofs

B.2.1. Proof of Corollary 3

Proof. Consider first the case where $T\mathcal{L}^{\mathsf{perish}} \lesssim L_T^{-1}$. Then:

$$L_T^{-1} + \sqrt{T\mathcal{L}^{\mathsf{perish}} L_T^{-1}} \lesssim L_T^{-1} \lesssim \sqrt{T},$$

since $L_T \gtrsim 1/\sqrt{T}$ by assumption. Thus,

$$\Delta_{\mathit{efficiency}} \lesssim \min\left\{\sqrt{T}, L_T^{-1} + \sqrt{T\mathcal{L}^{\mathsf{perish}} L_T^{-1}}\right\} + T\mathcal{L}^{\mathsf{perish}} \lesssim L_T^{-1},$$

where again we've used the assumption that $T\mathcal{L}^{\mathsf{perish}} \lesssim L_T^{-1}$.

For the bound on Δ_{EF}, we use the facts that $L_T \gtrsim 1/\sqrt{T}$ and $\mathcal{L}^{\mathsf{perish}} \lesssim 1/\sqrt{T}$ to obtain:

$$\Delta_{\mathit{EF}} \lesssim \max\{L_T, \mathcal{L}^{\mathsf{perish}} + 1/\sqrt{T}\} \lesssim L_T.$$

Suppose now $T\mathcal{L}^{\mathsf{perish}} \gtrsim L_T^{-1}$. In this case:

$$L_T^{-1} + \sqrt{T\mathcal{L}^{\mathsf{perish}} L_T^{-1}} \lesssim T\mathcal{L}^{\mathsf{perish}}.$$

Using the fact that $\mathcal{L}^{\mathsf{perish}} \lesssim 1/\sqrt{T}$, we obtain:

$$\Delta_{\mathit{efficiency}} \lesssim \min\left\{\sqrt{T}, L_T^{-1} + \sqrt{T\mathcal{L}^{\mathsf{perish}} L_T^{-1}}\right\} + T\mathcal{L}^{\mathsf{perish}} \lesssim T\mathcal{L}^{\mathsf{perish}}.$$

For the bound on Δ_{EF}, we similarly have $\Delta_{\mathit{EF}} \lesssim L_T$, since $\mathcal{L}^{\mathsf{perish}} \lesssim 1/\sqrt{T} \lesssim L_T$, by assumption. □

B.3. Section 4.3 Omitted Proofs

B.3.1. Proof of Lemma 3

Proof. Fix $t' < t$. Recall, for all $\tau \in [T]$, $\rho_{\tau,\theta} \geq |N_{\tau,\theta} - \mathbb{E}[N_{\tau,\theta}]|$, which implies

$$N_{\tau,\theta} \in [\mathbb{E}[N_{\tau,\theta}] - \rho_{\tau,\theta}, \mathbb{E}[N_{\tau,\theta}] + \rho_{\tau,\theta}].$$

Thus, from a simple application of Hoeffding's inequality (Lemma EC.1):

$$\mathbb{P}(|N_{(t,t']} - \mathbb{E}\left[N_{(t,t']}\right]| \geq \epsilon) \leq 2\exp\left(-\frac{2\epsilon^2}{\sum_\theta \sum_{\tau\in(t,t']} 4\rho_{\tau,\theta}^2}\right) \quad \text{(EC.2)}$$

We now consider our desired bound.

$$\mathbb{P}(|N_{(t,t']} - \mathbb{E}\left[N_{(t,t']}\right]| \leq \epsilon \,\forall t,t') \geq 1 - \sum_{t,t'} \mathbb{P}(|N_{(t,t']} - \mathbb{E}\left[N_{(t,t']}\right]| \geq \epsilon)$$
$$\geq 1 - \sum_{t,t'} 2\exp\left(-\frac{2\epsilon^2}{\sum_\theta \sum_{\tau\in(t,t']} 4\rho_{\tau,\theta}^2}\right) \geq 1 - \sum_{t,t'} 2\exp\left(-\frac{\epsilon^2}{2|\Theta|\rho_{\max}^2(t'-t)}\right)$$

where the first inequality follows from a union bound, the second inequality by plugging in Hoeffding's bound (EC.2), and the third inequality by upper bounding $\rho_{\tau,\theta}$ by $\rho_{\max}$, for all $\tau \in (t,t']$.
Solving for ϵ such that $2\exp\left(-\frac{\epsilon^2}{2|\Theta|\rho_{\max}^2(t'-t)}\right) = \delta/T^2$, we obtain our result. □

B.3.2. Proof of Lemma 5

Proof. The final high-probability bound follows from straightforward algebra, putting Lemmas 3 and 4 together. Indeed, we have that:

$$\mathbb{P}(\mathcal{E}) = 1 - \mathbb{P}(\mathcal{E}^c) = 1 - \mathbb{P}((\mathcal{E}_N \cap \mathcal{E}_{\overline{P}})^c).$$

Moreover:

$$\mathbb{P}(\mathcal{E}_N \cap \mathcal{E}_{\overline{P}}) = \mathbb{P}(\mathcal{E}_{\overline{P}} \mid \mathcal{E}_N)\mathbb{P}(\mathcal{E}_N) \geq (1-\delta)^2 \geq 1 - 2\delta \implies \mathbb{P}((\mathcal{E}_N \cap \mathcal{E}_{\overline{P}})^c) \leq 2\delta.$$

Plugging this in above, we obtain $\mathbb{P}(\mathcal{E}) \geq 1 - 2\delta$. □

B.3.3. Proof of Lemma 7

Proof. Consider first the case in which the algorithm always allocates $\underline{X}$ (i.e., $t_0 = T$). Then, the inequality is trivially satisfied and it suffices to prove the lower bound for $t_0 < T$. We have:

$$\begin{aligned}
B_T^{alg} &= B_{t_0}^{alg} - N_{t_0}\underline{X} - \text{PUA}_{t_0}^{alg} - N_{(t_0,T)}\overline{X} - \text{PUA}_{>t_0}^{alg} \\
&< N_{t_0}\overline{X} + \overline{N}_{>t_0}\underline{X} + \overline{P}_{t_0} - N_{t_0}\underline{X} - \text{PUA}_{t_0}^{alg} - N_{(t_0,T)}\overline{X} - \text{PUA}_{>t_0}^{alg} \\
&= N_{t_0}L_T + (\overline{N}_{>t_0} - N_{>t_0})\underline{X} + N_T\underline{X} - N_{(t_0,T)}L_T + \overline{P}_{t_0} - \text{PUA}_{\geq t_0}^{alg},
\end{aligned} \quad \text{(EC.3)}$$

where the first inequality follows from the fact that $B_{t_0}^{alg} < N_{t_0}\overline{X} + \overline{N}_{>t_0}\underline{X} + \overline{P}_{t_0}$ since $X_{t_0}^{alg} = \underline{X}$, and the second inequality uses $\overline{X} = \underline{X} + L_T$ and re-arranges terms. Since $\overline{X}$ was allocated at T, $B_T^{alg} - N_T\overline{X} \geq 0$, which then implies that $B_T^{alg} - N_T\underline{X} \geq N_T L_T$. Plugging this fact into (EC.3) and re-arranging, we obtain:

$$N_{t_0}L_T + (\overline{N}_{>t_0} - N_{>t_0})\underline{X} - N_{>t_0}L_T + \overline{P}_{t_0} - \text{PUA}_{\geq t_0}^{alg} > 0. \quad \text{(EC.4)}$$

We now upper bound the left-hand side of (EC.4). Using the facts that $\text{PUA}_{\geq t_0}^{alg} \geq 0$, $\underline{X} \leq \beta_{avg}$ by construction, and $\overline{P}_{t_0} \leq \overline{P}_1 = \overline{\Delta}(\underline{X})$, we have, for $C = \sqrt{2|\Theta|\rho_{max}^2 \log(2T^2/\delta)}$:

$$0 < N_{t_0} L_T + (\overline{N}_{>t_0} - N_{>t_0})\underline{X} - N_{>t_0} L_T + \overline{P}_{t_0} - \text{PUA}_{\geq t_0}^{alg} \leq \rho_{\max} L_T + 2C\beta_{avg}\sqrt{T - t_0} - L_T(T - t_0) + \overline{\Delta}(\underline{X}). \tag{EC.5}$$

Consider now the quadratic function $f(x) = -L_T x^2 + 2C\beta_{avg} x + \rho_{\max} L_T + \overline{\Delta}(\underline{X})$, which has a positive root at:

$$\begin{aligned} x^+ &= \frac{2C\beta_{avg} + \sqrt{4C^2\beta_{avg}^2 + 4L_T\left(\rho_{\max} L_T + \overline{\Delta}(\underline{X})\right)}}{2L_T} \\ &= \frac{C\beta_{avg}}{L_T} + \sqrt{\frac{C^2\beta_{avg}^2}{L_T^2} + \frac{\rho_{\max} L_T + \overline{\Delta}(\underline{X})}{L_T}} \\ &\leq 2\frac{C\beta_{avg}}{L_T} + \sqrt{\frac{\overline{\Delta}(\underline{X})}{L_T}} + \sqrt{\rho_{\max}} \\ &< c\left(\frac{1}{L_T} + \sqrt{\frac{\overline{\Delta}(\underline{X})}{L_T}}\right), \end{aligned}$$

for some $c \in \widetilde{\Theta}(1)$. Thus, for all $x \geq c\left(\frac{1}{L_T} + \sqrt{\frac{\overline{\Delta}(\underline{X})}{L_T}}\right)$, $f(x) \leq 0$. Letting $x = \sqrt{T - t_0}$, we obtain that the right-hand side of (EC.5) is non-positive for all t_0 such that $T - t_0 \geq c^2\left(\frac{1}{L_T} + \sqrt{\frac{\overline{\Delta}(\underline{X})}{L_T}}\right)^2 \iff t_0 \leq T - c^2\left(\frac{1}{L_T} + \sqrt{\frac{\overline{\Delta}(\underline{X})}{L_T}}\right)^2$, which would lead to a contradiction.

Concluding, we have $t_0 > T - c^2\left(\frac{1}{L_T} + \sqrt{\frac{\overline{\Delta}(\underline{X})}{L_T}}\right)^2 \geq T - \tilde{c}^2\left(\frac{1}{L_T} + \sqrt{\frac{T\mathcal{L}^{\text{perish}}}{L_T}}\right)^2$, where the final inequality follows from the fact that $\overline{\Delta}(\underline{X}) \leq B - \overline{N}\underline{X} = \overline{N}\mathcal{L}^{\text{perish}} \lesssim T\mathcal{L}^{\text{perish}}$. □

B.3.4. Proof of Lemma 8

Proof. We show the two properties by induction on t.

Base Case $t = 1$. By definition, $\mathcal{B}_1^{alg} = \mathcal{B} = \mathcal{B}_t(\underline{X})$. We now argue that $\text{PUA}_1^{alg} \leq \text{PUA}_1(\underline{X})$. Suppose there exists a resource b which perished at the end of $t = 1$. Then, either:

1. b was neither allocated by our algorithm, nor under the $\underline{X}$ allocation. Hence, it perished unallocated under both allocations.
2. b was allocated by our algorithm but not by the $\underline{X}$ allocation. Hence, it perished unallocated under $\underline{X}$ but not our algorithm.
3. b was allocated under the $\underline{X}$ allocation but not by our algorithm. This could never hold, since both algorithms begin with the same set of resources and our algorithm allocated (weakly) more than $\underline{X}$, and under the same ordering σ.

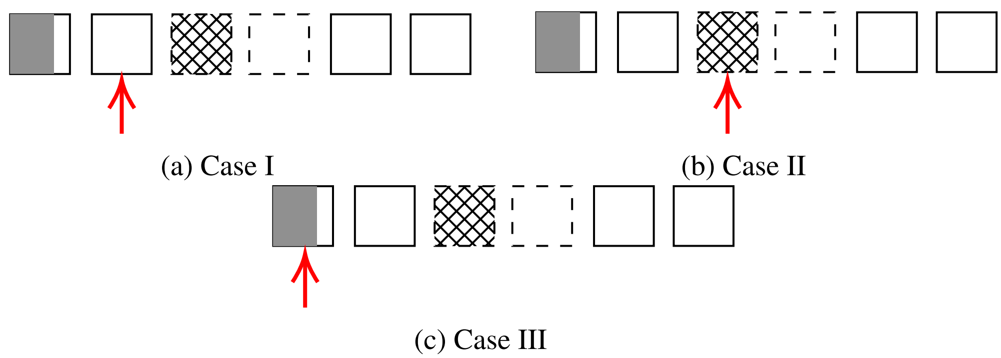

Figure EC.1 Illustration of the three cases in the induction step of the proof of the fact that $\text{PUA}^{alg}_{t+1} \leq \text{PUA}_{t+1}(\underline{X})$ (Lemma 8). Here, we assume $\mathcal{B}^{alg}_{t+1} \subseteq \mathcal{B}_{t+1}(\underline{X})$. The squares across all three plots show the resources in $\mathcal{B}_{t+1}(\underline{X})$, ordered left to right according to σ; the dashed squares correspond to $\mathcal{B}^{alg}_{t+1}$. The gray-shaded region corresponds to resource fractionally allocated by the $\underline{X}$ process at the beginning of $t+1$, and the cross-hatched region to the set of resources allocated by our algorithm. Finally, the red arrow corresponds to the resource b considered in each case.

Step case $t \to t+1$. We first show that $\mathcal{B}^{alg}_{t+1} \subseteq \mathcal{B}_{t+1}(\underline{X})$. Let $b \in \mathcal{B}^{alg}_{t+1}$. Then b did not perish, and moreover $b \in \mathcal{B}^{alg}_{t}$. Then, by the inductive hypothesis, $b \in \mathcal{B}_t(\underline{X})$. Consider the following cases:

1. b was not allocated under the $\underline{X}$ process. Then, $b \in \mathcal{B}_{t+1}(\underline{X})$, since it did not perish.
2. b was allocated under the $\underline{X}$ process. In this case, since the algorithm allocated (weakly) more than $\underline{X}$ according to the same ordering, this resource *must* have been available to both the algorithm and the $\underline{X}$ process. This then contradicts that b was not allocated by the algorithm.

We now argue that $\text{PUA}^{alg}_{t+1} \leq \text{PUA}_{t+1}(\underline{X})$. Suppose there exists a resource b that perished at time $t+1$. We consider the following cases (see Figure EC.1 for an illustration):

1. b was neither allocated by our algorithm, nor under the $\underline{X}$ allocation. Hence, it perished unallocated under both allocations.
2. b was allocated by our algorithm but not by the $\underline{X}$ allocation. Hence, it perished unallocated under $\underline{X}$ but not our algorithm.
3. b was allocated under the $\underline{X}$ allocation but not by our algorithm. Then, b *must* have either perished or been allocated before $t+1$, since the set of remaining resources under our algorithm is (weakly) nested in the set of remaining resources under the $\underline{X}$ for all $t' \leq t$, by the inductive hypothesis. Thus, b could not have perished at the end of $t+1$ under our algorithm's sample path.

□

B.4. Section 4.4 Omitted Proofs

For ease of notation, we let $\nu_t = \mathbb{E}[P_{<t}]$ for all $t \in \{2,\ldots,T\}$.

B.4.1. Proof of Proposition 1

Proof. Let $t \in [T]$ be such that $\mathbb{E}[P_{<t}] > t-1$. Then:

$$\mathbb{P}(P_{<t} \le t-1\ \forall\, t \ge 2) \le \mathbb{P}(P_{<t} \le t-1) = \mathbb{P}\left(\sum_{b\in\mathcal{B}} \mathbb{1}\{T_b < t\} \le t-1\right), \tag{EC.6}$$

where the second equality is by definition.

Consider first the case where $\mathcal{B}^{rand}_{<t} = \emptyset$. In this case, if b perishes before t with strictly positive probability, it must be that $b \in \mathcal{B}^{det}_{<t}$. Then:

$$\mathbb{P}\left(\sum_{b\in\mathcal{B}} \mathbb{1}\{T_b < t\} \le t-1\right) = \mathbb{P}\left(\sum_{b\in\mathcal{B}^{det}_{<t}} \mathbb{1}\{T_b < t\} \le t-1\right) = \mathbb{P}\left(|\mathcal{B}^{det}_{<t}| \le t-1\right), \tag{EC.7}$$

where the second equality follows from the fact that items in $\mathcal{B}^{det}_{<t}$ perish before t with probability 1. By the same reasoning:

$$t-1 < \mathbb{E}[P_{<t}] = \sum_{b\in\mathcal{B}} \mathbb{P}(T_b < t) = \sum_{b\in\mathcal{B}^{det}_{<t}} \mathbb{P}(T_b < t) = |\mathcal{B}^{det}_{<t}| \implies \mathbb{P}(|\mathcal{B}^{det}_{<t}| \le t-1) = 0.$$

Plugging this back into (EC.6), we obtain $\mathbb{P}(P_{<t} \le t-1\ \forall\, t \ge 2) = 0$.

Consider now the case where $\mathcal{B}^{rand}_{<t} \neq \emptyset$. The goal is to show the existence of ϵ such that $\mathbb{P}(P_{<t} \le t-1\ \forall t \ge 2) \le \epsilon$. Define the random variable:

$$Y_b = \mathbb{1}\{T_b < t\} - \mathbb{P}(T_b < t), \quad b \in \mathcal{B}^{rand}_{<t}.$$

By construction, $\mathbb{E}[Y_b] = 0, 0 < \mathbb{E}[Y_b^2] \le 1$, and $\mathbb{E}[|Y_b|^3] \le 1$. We have:

$$\begin{aligned}\mathbb{P}(P_{<t} \le t-1) &= \mathbb{P}\left(\sum_{b\in\mathcal{B}^{rand}_{<t}} \mathbb{1}\{T_b < t\} \le t-1-|\mathcal{B}^{det}_{<t}|\right)\\ &= \mathbb{P}\left(\sum_{b\in\mathcal{B}^{rand}_{<t}} Y_b \le t-1-|\mathcal{B}^{det}_{<t}| - \sum_{b\in\mathcal{B}^{rand}_{<t}} \mathbb{P}(T_b < t)\right).\end{aligned}$$

By assumption, $\mathbb{E}[P_{<t}] = |\mathcal{B}^{det}_{<t}| + \sum_{b\in\mathcal{B}^{rand}_{<t}} \mathbb{P}(T_b < t) > t-1$. Hence,

$$\mathbb{P}(P_{<t} \le t-1) \le \mathbb{P}\left(\sum_{b\in\mathcal{B}^{rand}_{<t}} Y_b \le 0\right) = \mathbb{P}\left(\frac{\sum_{b\in\mathcal{B}^{rand}_{<t}} Y_b}{\sqrt{\sum_{b\in\mathcal{B}^{rand}_{<t}} \mathbb{E}[Y_b^2]}} \le 0\right).$$

Let $\Phi(\cdot)$ denote the cdf of the standard normal distribution. By the Berry-Esseen Theorem,

$$\mathbb{P}\left(\frac{\sum_{b\in\mathcal{B}^{rand}_{<t}} Y_b}{\sqrt{\sum_{b\in\mathcal{B}^{rand}_{<t}} \mathbb{E}[Y_b^2]}} \le 0\right) \le \Phi(0) + \left(\sum_{b\in\mathcal{B}^{rand}_{<t}} \mathbb{E}[Y_b^2]\right)^{-3/2} \cdot \sum_{b\in\mathcal{B}^{rand}_{<t}} \mathbb{E}[|Y_b|^3]$$

$$
\begin{aligned}
&= \frac{1}{2} + \left(\sum_{b \in \mathcal{B}_{<t}^{rand}} \text{Var}[\mathbb{1}\{T_b < t\}] \right)^{-3/2} \cdot \sum_{b \in \mathcal{B}_{<t}^{rand}} \mathbb{E}[|Y_b|^3] \\
&= \frac{1}{2} + \left(\text{Var}\left[\sum_{b \in \mathcal{B}_{<t}^{rand}} \mathbb{1}\{T_b < t\} \right] \right)^{-3/2} \cdot \sum_{b \in \mathcal{B}_{<t}^{rand}} \mathbb{E}[|Y_b|^3] \\
&\leq \frac{1}{2} + \text{Var}[P_{<t}]^{-3/2} \cdot T \\
&= \frac{1}{2} + \text{Std}\left[P_{<t}\right]^{-3} \cdot T.
\end{aligned}
$$

Putting this all together, we obtain:

$$
\mathbb{P}(P_{<t} \leq t-1 \ \forall t \geq 2) \leq \frac{1}{2} + \text{Std}\left[P_{<t}\right]^{-3} \cdot T.
$$

As a result, the perishing process cannot be δ-offset-expiring for any $\delta < \frac{1}{2} - \text{Std}\left[P_{<t}\right]^{-3} \cdot T$. □

B.4.2. Proof of Proposition 2

Proof. Recall, $P_{<t} = |\mathcal{B}_{<t}^{det}| + \sum_{b \in \mathcal{B}_{<t}^{rand}} \mathbb{1}\{T_b < t\}$. By Chebyshev's inequality:

$$
\mathbb{P}(P_{<t} > t-1) = \mathbb{P}\left(P_{<t} - \nu_t \geq t - \nu_t\right) \leq \left(\frac{\text{Std}\left[P_{<t}\right]}{t - \nu_t} \right)^2.
$$

Similarly, by Hoeffding's inequality we have:

$$
\mathbb{P}(P_{<t} > t-1) = \mathbb{P}\left(\sum_{b \in \mathcal{B}_{<t}^{rand}} \mathbb{1}\{T_b < t\} - \mathbb{P}(T_b < t) \geq t - \nu_t \right) \leq \exp\left(- \frac{2\left(t - \nu_t\right)^2}{|\mathcal{B}_{<t}^{rand}|} \right).
$$

Then, via a straightforward union bound we have:

$$
\mathbb{P}(P_{<t} \leq t-1 \ \forall t \geq 2) \geq 1 - \sum_{t=2}^{T} \mathbb{P}(P_{<t} > t-1) \geq 1 - \sum_{t=2}^{T} \min\left\{ \left(\frac{\text{Std}\left[P_{<t}\right]}{t - \nu_t} \right)^2, \exp\left(- \frac{2(t-\nu_t)^2}{|\mathcal{B}_{<t}^{rand}|} \right) \right\}.
$$

□

B.4.3. Proof of Proposition 3

Proof.

For all $b \in \mathcal{B}$, we have:

$$
\mathbb{P}(T_b \leq t-1) = 1 - (1-p)^{t-1} \implies \begin{cases} \nu_t &= T\left(1 - (1-p)^{t-1}\right) \\ \text{Var}[P_{<t}] &= T\left(1 - (1-p)^{t-1}\right)(1-p)^{t-1}. \end{cases} \tag{EC.8}
$$

By Proposition 2, δ-offset-expiry holds for all $\delta \geq \sum_{t=2}^{T} \left(\frac{\text{Std}[P_{<t}]}{t - \nu_t} \right)^2$. By (EC.8):

$$
\left(\frac{\text{Std}\left[P_{<t}\right]}{t - \nu_t} \right)^2 = \frac{T\left(1 - (1-p)^{t-1}\right)(1-p)^{t-1}}{\left(t - T\left(1 - (1-p)^{t-1}\right)\right)^2}
$$

Using the fact that $(1-p)^{t-1} \geq 1-(t-1)p$, we have:

$$t - T(1-(1-p)^{t-1}) \geq t - T(t-1)p > 0,$$

where the final inequality follows from the assumption that $p \leq 1/T$. Taking derivatives, it is easy to show that the function $f(x) = \frac{(1-x)x}{(t-T(1-x))^2}$ is decreasing for $t - T(1-x) > 0$. Hence, leveraging the same lower bound on $(1-p)^{t-1}$ we have:

$$\begin{aligned}
\left(\frac{\text{Std}\left[P_{<t}\right]}{t-\nu_t}\right)^2 &\leq \frac{T(t-1)p(1-(t-1)p)}{(t-T(t-1)p)^2} \leq \frac{Ttp(1-(t-1)p)}{t^2\,(1-Tp)^2} \leq \frac{Tp(1-(t-1)p)}{(t-1)(1-Tp)^2} \\
\implies \sum_{t=2}^{T}\left(\frac{\text{Std}\left[P_{<t}\right]}{t-\nu_t}\right)^2 &\leq \frac{Tp}{(1-Tp)^2}\sum_{t=2}^{T}\left(\frac{1}{t-1}-p\right) \\
&= \frac{Tp}{(1-Tp)^2}\left(\sum_{t=1}^{T-1}\frac{1}{t} - p(T-1)\right) \\
&\leq \frac{Tp}{(1-Tp)^2}\cdot 2\log T.
\end{aligned}$$

Hence, the perishing process if δ-offset-expiring for all $\delta \geq 2\log T \cdot \frac{Tp}{(1-Tp)^2}$.

We conclude by showing the lower bound on $\underline{X}$. By definition,

$$\begin{aligned}
\overline{\Delta}(X) &= \mu(X) + \frac{1}{2}\left(\log(3\log(T)/\delta) + \sqrt{\log^2(3\log(T)/\delta) + 8\mu(X)\log(3\log(T)/\delta)}\right) \\
&\leq \mu(X) + \frac{1}{2}\left(2\log(3\log(T)/\delta) + \frac{1}{2\log(3\log(T)/\delta)}\cdot 8\mu(X)\log(3\log(T)/\delta)\right) \\
&= 3\mu(X) + \log(3\log(T)/\delta),
\end{aligned}$$

where the inequality follows from concavity. Moreover:

$$\begin{aligned}
\mu(X) = \sum_b \mathbb{P}(T_b < \min\{T, \tau_b(1 \mid X, \sigma)\}) \leq \sum_b \mathbb{P}(T_b < T) &= T(1-(1-p)^{T-1}) \\
\implies \overline{\Delta}(X) &\leq 3T(1-(1-p)^{T-1}) + \log(3\log(T)/\delta) \\
&\leq 3T(T-1)p + \log(3\log(T)/\delta).
\end{aligned}$$

Since any feasible stationary allocation X must satisfy $X \leq \frac{T-\overline{\Delta}(X)}{T}$, it suffices to have:

$$X \leq \frac{T - 3T(T-1)p - \log(3\log(T)/\delta)}{T} = 1 - 3(T-1)p - \frac{\log(3\log(T)/\delta)}{T}. \tag{EC.9}$$

Noting that the right-hand side of (EC.9) is non-negative for $\delta \geq 3\log T \cdot \exp(-(T-3T^2p))$, and that $2\log T \cdot \frac{Tp}{(1-Tp)^2} \geq 3\log T \cdot \exp(-(T-3T^2p))$ for $p = o(1)$, we obtain the result. $\square$

B.4.4. Proof of Proposition 4

Proof. For ease of notation, we let $\mu_b = \mathbb{E}[T_b]$. For a stationary allocation $X = 1 - T^{-\alpha}$, let $\mu(X) = \sum_b \mathbb{P}(T_b < \min\{T, \lceil \frac{\sigma(b)}{X} \rceil\})$, and $\mu_b = \mathbb{E}[T_b]$. By Chebyshev's inequality, we have:

$$\mu(X) \leq \sum_b \mathbb{P}\left(T_b - \mu_b \leq \min\left\{T, \lceil \frac{\sigma(b)}{X} \rceil\right\} - \mu_b\right) \leq \sum_b \mathbb{P}\left(|T_b - \mu_b| \geq \mu_b - \min\left\{T, \lceil \frac{\sigma(b)}{X} \rceil\right\}\right)$$
$$\leq \sum_b \frac{\mathrm{Var}[T_b]}{\left(\mu_b - \min\{T, \lceil \frac{\sigma(b)}{X} \rceil\}\right)^2}, \tag{EC.10}$$

where we used the assumption that $\mu_b > \min\{T, \lceil \frac{\sigma(b)}{X} \rceil\}$ for all $b \in \mathcal{B}$. However, by assumption we have

$$\mu(X) \leq \sum_b \frac{\mathrm{Var}[T_b]}{\left(\mu_b - \min\{T, \lceil \frac{\sigma(b)}{X} \rceil\}\right)^2} \leq \frac{1}{2} T^{1-\alpha}$$

Moreover, by definition:

$$\overline{\Delta}(X) = \mu(X) + \frac{1}{2}\left(\log(3\log(T)/\delta) + \sqrt{\log^2(3\log(T)/\delta) + 8\mu(X)\log(3\log(T)/\delta)}\right)$$
$$\leq \mu(X) + \log(3\log(T)/\delta) + \sqrt{2\mu(X)\log(3\log(T)/\delta)}.$$

Since any feasible X must satisfy $X \leq 1 - \overline{\Delta}(X)/T$, we have that $\overline{\Delta}(X) \leq T^{1-\alpha}$ for $X = 1 - T^{-\alpha}$. Thus, it suffices for $\mu(X)$ to satisfy

$$\mu(X) + \sqrt{2\mu(X)\log(3\log(T)/\delta)} + \log(3\log(T)/\delta) \leq T^{1-\alpha}$$

We have that $\mu(X) \leq \frac{1}{2} T^{1-\alpha}$ satisfies this inequality for all $\delta \geq 3\log(T) e^{-\frac{1}{8}T^{1-\alpha}}$. □

Appendix C: Simulation details

Computing Infrastructure. The experiments were conducted on a personal computer with an Apple M2, 8-core processor and 16.0GB of RAM.

Real-World Simulations. We use the dataset of Keskin et al. (2022), which contains detailed daily information on orders, inventory, sales, prices, costs, and holiday indicators for a variety of perishable fruits and vegetables across multiple stores, for a leading supermarket chain in China. Following Keskin et al. (2022), we focus on "ginger" in store "A," as this product experiences no stock-outs over the sample period. The data span the full calendar year of 2013 and, for each day, include beginning inventory, replenishment quantity, realized sales, ending inventory, price, and a holiday indicator. All inventory and sales quantities are normalized in the dataset, so there is no explicit notion of physical "units." Throughout, we treat these normalized quantities as inventory units for modeling purposes.

As in Keskin et al. (2022), we assume demand is normally distributed and fit the mean and standard deviation using the full year of data. Specifically, we treat observed daily sales as a proxy for demand

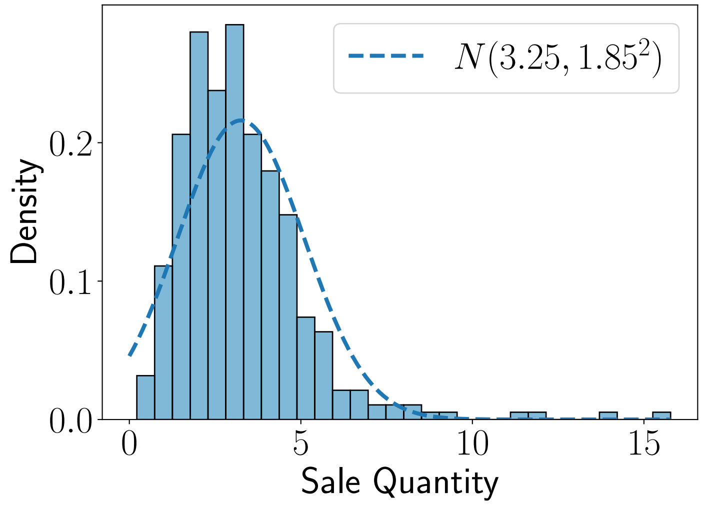

Figure EC.2 Empirical distribution of daily ginger sales with fitted normal distribution using the sample mean and standard deviation.

and fit a normal distribution to this data, obtaining mean $\mu = 3.2$ and standard deviation $\sigma = 1.85$.[4] See Figure EC.2 for a histogram of the true sales data against the normal approximation.

To estimate the perishability rate, we fit a maximum likelihood estimator under the assumption that each unit independently perishes each day with probability p. We compute the number of units that perish on each day t as follows:

$$\text{Perish}_t = \text{BeginStockQty}_t + \text{Restock}_t - \text{SalesQty}_t - \text{EndStockQty}_t,$$

where Restock_t is used to denote the replenishment quantity on day t. (The dataset reports the replenishment quantity separately from the initial inventory level for each day t.) Under the assumption that each unit perishes i.i.d. with probability p on each day, the number of perished units on day t is binomially distributed with BeginStockQty_t trials and parameter p. Hence, the maximum likelihood estimate of p is given by:

$$\hat{p} = \frac{\sum_t \text{Perish}_t}{\sum_t \text{BeginStockQty}_t}.$$

On the ginger dataset, this results in $\hat{p} = 0.0024$.

[4] This specification ignores potential effects from price changes, seasonality, or holiday-related demand shifts, as these effects are not the focus of this work.

C.1. Additional Results

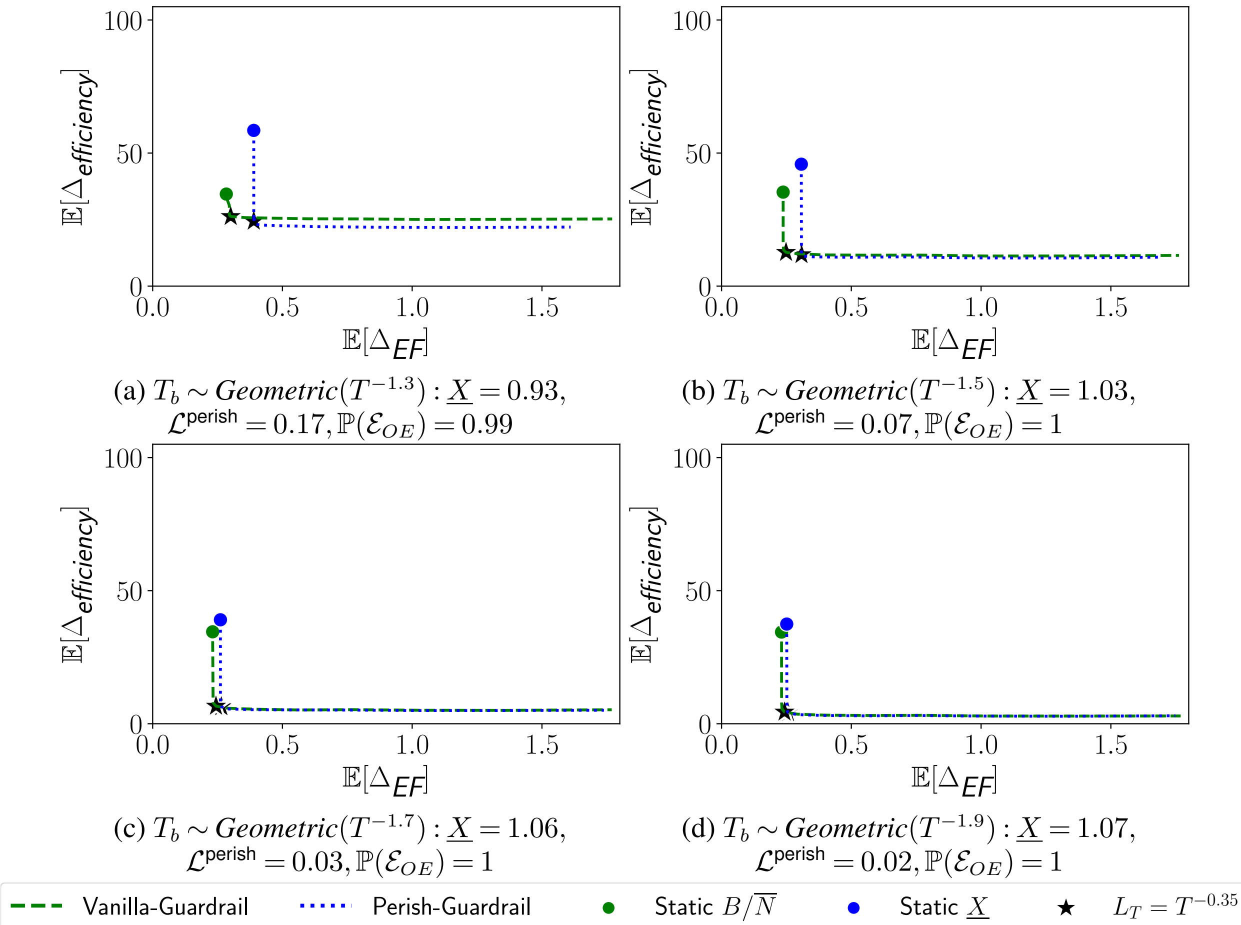

(a) $T_b \sim \textit{Geometric}(T^{-1.3}) : \underline{X} = 0.93$, $\mathcal{L}^{\text{perish}} = 0.17, \mathbb{P}(\mathcal{E}_{OE}) = 0.99$

(b) $T_b \sim \textit{Geometric}(T^{-1.5}) : \underline{X} = 1.03$, $\mathcal{L}^{\text{perish}} = 0.07, \mathbb{P}(\mathcal{E}_{OE}) = 1$

(c) $T_b \sim \textit{Geometric}(T^{-1.7}) : \underline{X} = 1.06$, $\mathcal{L}^{\text{perish}} = 0.03, \mathbb{P}(\mathcal{E}_{OE}) = 1$

(d) $T_b \sim \textit{Geometric}(T^{-1.9}) : \underline{X} = 1.07$, $\mathcal{L}^{\text{perish}} = 0.02, \mathbb{P}(\mathcal{E}_{OE}) = 1$

Figure EC.3 Empirical trade-off between $\Delta_{efficiency}$ and Δ_{EF} for the different algorithms under various values of L_T.

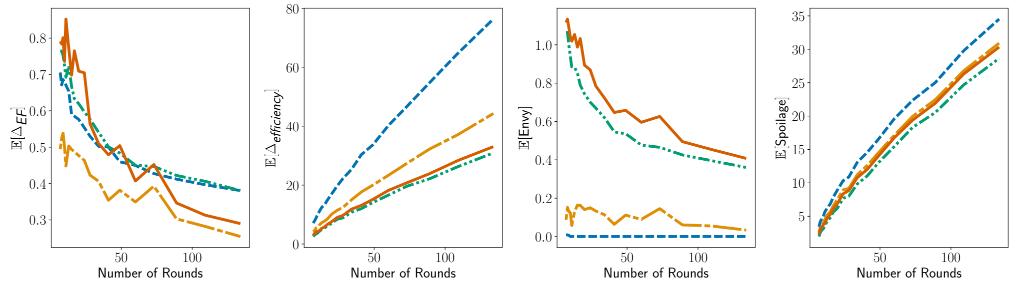

(a) $T_b \sim \textit{Geometric}(T^{-(1.3)})$

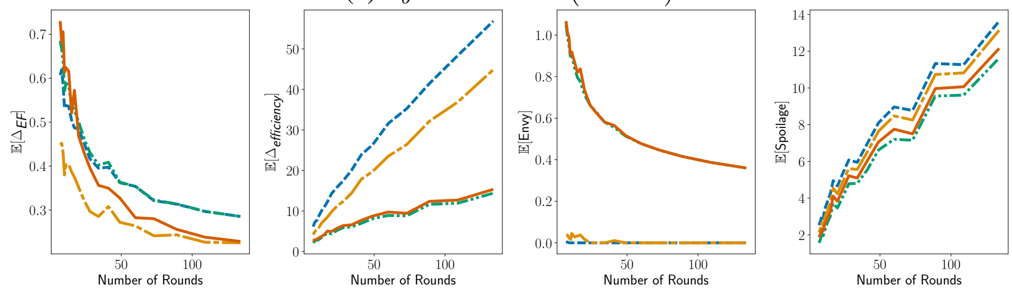

(b) $T_b \sim \textit{Geometric}(T^{-(1.5)})$

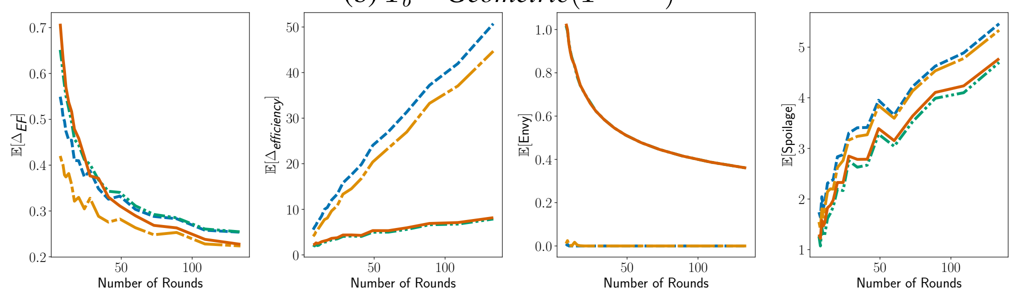

(c) $T_b \sim \textit{Geometric}(T^{-(1.7)})$

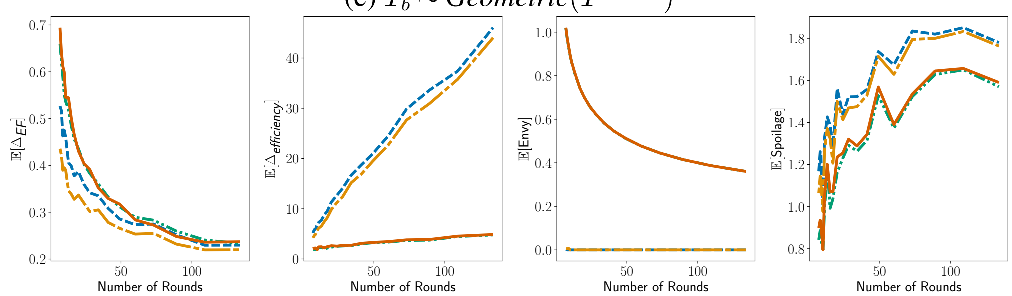

(d) $T_b \sim \textit{Geometric}(T^{-(1.9)})$

Static $\underline{X}$ — Static $B/\overline{N}$ — Perish-Guardrail — Vanilla-Guardrail

Figure EC.4 Numerical results for $T_b \sim \textit{Geometric}(p)$ as described in Section 5.2.

Appendix D: Useful lemmas

We use the following standard theorems throughout the proof. See, e.g. Vershynin (2018) for proofs and further discussion.

LEMMA EC.1 **(Hoeffding's Inequality (Vershynin 2018)).** *Let $X_1, \ldots, X_n$ be independent random variables such that $a_i \le X_i \le b_i$ almost surely, with $S_n = \sum_i X_i$. Then, for all $t > 0$:*

$$\mathbb{P}(|S_n - \mathbb{E}[S_n]| \ge t) \le 2\exp\left(-\frac{2t^2}{\sum_i (b_i - a_i)^2}\right).$$

The next is a Chernoff bound for Bernoulli random variables. See Mitzenmacher and Upfal (2017).

LEMMA EC.2 **(Chernoff Bound for Sum of Bernoulli Random Variables).** *Consider a sequence of Bernoulli random variables $(X_i)_{i\in[N]}$, independently distributed with probability of success $p_i \in (0,1)$. Let $X = \sum_i X_i$, and let $\mu = \mathbb{E}[X] = \sum_i p_i$. Then, for all $\epsilon > 0$:*

$$\mathbb{P}(X \ge (1+\epsilon)\mu) \le \exp\left(-\frac{\epsilon^2}{2+\epsilon}\mu\right).$$

COROLLARY EC.1. *Consider a sequence of Bernoulli random variables $(X_i)_{i\in[N]}$, independently distributed with probability of success $p_i \in (0,1)$. Let $X = \sum_i X_i$, and let $\mu = \mathbb{E}[X] = \sum_i p_i$. Then, for any $\delta > 0$, with probability at least $1-\delta$ we have:*

$$X \le \mu + \frac{1}{2}\left(\log(1/\delta) + \sqrt{\log^2(1/\delta) + 8\mu\log(1/\delta)}\right).$$

Proof. Setting the right hand side equal to δ in Lemma EC.2 and solving for ϵ, we have:

$$\frac{\epsilon^2}{2+\epsilon}\mu = \log(1/\delta) \iff \epsilon^2\mu - \epsilon\log(1/\delta) - 2\log(1/\delta) = 0$$
$$\iff \epsilon = \frac{\log(1/\delta) + \sqrt{\log^2(1/\delta) + 8\mu\log(1/\delta)}}{2\mu}.$$

Plugging this value of ϵ into $(1+\epsilon)\mu$ we have the result. □