Appendix A: Ommitted Proofs

A.1. Preliminaries

Proof of Lemma 1. The Lemma follows from Hoeffding's Inequality (Hoeffding 1994), which asserts that for independent random variables $X_1, \ldots, X_n$, where $X_i \in \{0,1\} \forall i$, we have

$$\mathbb{P}\left[|\sum_{i=1}^{n} X_i - \mathbb{E}[\sum_{i=1}^{n} X_i]| \geq \psi \right] \leq 2e^{-\frac{2\psi^2}{n}}.$$

With $n = \bar{t}$, X_i denoting the indicator of whether the ith remaining arrival is of type θ, and $\psi = \sqrt{\alpha \bar{t} \log(\bar{t})}$:

$$\mathbb{P}\left[|\bar{N}(t) - \bar{t} p_{\theta|} \geq \sqrt{\alpha \bar{t} \log(\bar{t})} \right] \leq 2e^{-\frac{2\sqrt{\alpha \bar{t} \log(\bar{t})}^2}{\bar{t}}} = 2\bar{t}^{-2\alpha}.$$

The result follows by taking a union bound over all k types θ. □

Proof of Proposition 1. The proof of part (ii) of the proposition heavily relies on the following result from the literature (where the variable names are adapted to our notation):

PROPOSITION 5 **(Theorem 2.4 in Mangasarian and Shiau (1987)).** *Let the linear program*

$$\min_{\vec{x}, \vec{z}} \vec{c_1} \cdot \vec{x} + \vec{c_2} \cdot \vec{z} \quad s.t.\ C_1 \begin{bmatrix} \vec{x} \\ \vec{z} \end{bmatrix} \leq \vec{M},\ C_2 \vec{x} = \vec{N}.$$

have non-empty optimal solution sets S^1 and S^2 for right-hand sides $(\vec{M}^1, \vec{N}^1)$ and $(\vec{M}^2, \vec{N}^2)$ respectively. For each $(\vec{x}^1, \vec{z}^1) \in S^1$ there exists $(\vec{x}^2, \vec{z}^2) \in S^2$ such that $|\vec{x}^1 - \vec{x}^2|_\infty \leq v(C_1; C_2) \left\| \begin{bmatrix} N^1 - N^2 \\ M^1 - M^2 \end{bmatrix} \right\|_\infty$ where

$$v(C_1, C_2) = \sup_{\vec{u}, \vec{v}} \left\{ \max\{\|u\|_1, \|v\|_1\} : \|uC_1 + vC_2\|_1 = 1, \textit{rows of } \begin{pmatrix} C_1 \\ C_2 \end{pmatrix} \textit{ corresponding to nonzero elements of } \begin{pmatrix} u \\ v \end{pmatrix} \textit{ are lin. ind.} \right\}$$

To adopt the result for our purposes, set $\vec{M}^1 = \begin{bmatrix} -\vec{y} \\ \vec{B}|\vec{N}^1|_1 \end{bmatrix}$, $\vec{M}^2 = \begin{bmatrix} -\vec{y} \\ \vec{B}|\vec{N}^2|_1 \end{bmatrix}$, and let C_1 and C_2 be such that $C_1 \begin{bmatrix} \vec{x} \\ \vec{z} \end{bmatrix} \leq \vec{M}^1$ encodes the constraints $-x_{\theta j} \leq -y_{\theta j}$, and $A\vec{x} - Dz \leq \vec{B}|\vec{N}^1|_1$ and C_2 encodes the constraints $\sum_j x_{\theta j} = N_\theta^1$, respectively $A\vec{x} - Dz \leq \vec{B}|\vec{N}^2|_1$ and $\sum_j x_{\theta j} = N_\theta^2$ for $(\vec{M}^2, \vec{N}^2)$. Then the proposition says that there exists δ, dependent only on C_1 and C_2 (which are independent of T, $B, \vec{c}_1, \vec{c}_2$) such that for any $(\vec{x}^1, \vec{z}^1)$ optimal for $\vec{N}^1$ we have some $(\vec{x}^2, \vec{z}^2)$ optimal for $\vec{N}^2$ where

$$\left\| \begin{bmatrix} \vec{x}^1 - \vec{x}^2 \\ \vec{z}^1 - \vec{z}^2 \end{bmatrix} \right\|_\infty \leq \delta \left\| \begin{bmatrix} \vec{N}_1 - \vec{N}_2 \\ \vec{M}_1 - \vec{M}_2 \end{bmatrix} \right\|_\infty = \delta \left\| \begin{bmatrix} \vec{N}_1 - \vec{N}_2 \\ 0 \end{bmatrix} \right\|_\infty,$$ so in particular $|\vec{x}_1 - \vec{x}_2|_\infty \leq \delta |\vec{N}_1 - \vec{N}_2|_\infty$.

The proof of part (i) follows as the inner objective is Lipschitz in $\vec{x}$ and $\vec{z}$, and the corresponding optimal solutions $\vec{x}^1, \vec{x}^2$ and $\vec{z}^1, \vec{z}^2$ are Lipschitz in $\vec{N}$ (from part ii), so $f_{\vec{N}}(\vec{x})$ is Lipschitz in $\vec{x}$. $\square$

A.2. Proofs of Lemmas in Section 3

Proof of Lemma 2. Observe from $\texttt{OPT}^{\,t}$ that the optimization problems $\texttt{OPT}^{t_a}$ and $\texttt{OPT}^{t_b}$ vary only with respect to the constraints, where $\vec{x}^{\texttt{ALG}}(t_a-1)$ is replaced by $\vec{x}^{\texttt{ALG}}(t_b-1)$. By the lemma assumption, the solution to $\texttt{OPT}^{t_a}$ fulfills that $x^{\texttt{ALG}}_{\theta j}(t_b-1) \le \texttt{OPT}^{t_a}_{\theta j}$. Thus, it is feasible to $\texttt{OPT}^{t_b}$, implying that the objective of $\texttt{OPT}^{t_b} \ge \texttt{OPT}^{t_a}$. Equality follows as $\texttt{OPT}^{t_b}$ is more constrained than $\texttt{OPT}^{t_a}$. $\square$

Proof of Lemma 3. Observe first that, by Proposition 1, we have

$$\max_{\theta,j} |\texttt{OPT}_{\theta j} - T\times\texttt{DLP}_{\theta j}| \le \delta \max_{\theta} |Tp_\theta - N_\theta|, \quad \text{implying that } \forall \theta, j : \texttt{OPT}_{\theta j} \ge T\times\texttt{DLP}_{\theta j} - \delta \max_{\theta} |Tp_\theta - N_\theta|.$$

Setting $\alpha = \frac{(T\log(T))^\epsilon}{\delta^2}$ and $t=0$, i.e., $\bar{t}=T$, in Lemma 1 we find with probability greater $1-2kT^{-\frac{T^\epsilon}{\delta^2}}$ that the expression $\delta \max_\theta |Tp_\theta - N_\theta|$ is at most $[T\log(T)]^{\frac{1+\epsilon}{2}}$, which implies the result. $\square$

Proof of Lemma 4. Observe that the first t arrivals, $N_\theta(t)$ and the last t arrivals, $\bar{N}_\theta(t)$ come from the same distribution; thus, Lemma 1 gives $\mathbb{P}\left[\max_\theta\{|N_\theta(t) - tp_\theta|\} > \sqrt{\alpha t \log(t)}\right] \le 2kt^{-2\alpha}$. With $\alpha = \frac{(t\log(t))^\epsilon}{\delta^2}$, we get

$$\mathbb{P}[\mathcal{S}_t] \le \mathbb{P}\left[\max_\theta\{|N_\theta(t) - tp_\theta|\} > \frac{[t\log(t)]^{\frac{1+\epsilon}{2}}}{\delta}\right] \le 2kt^{-2\frac{(t\log(t))^\epsilon}{\delta^2}} < 2kt^{\frac{-2t^\epsilon}{\delta^2}}. \quad \square$$

Proof of Lemma 5. We distinguish between periods t where the bad event $\mathcal{S}_t$ holds true, and ones where it does not. The former we charge to the right-hand side; this leaves us with showing that across the periods in which $\mathcal{S}_t$ does not hold true, with probability at least $1-\frac{1}{T}$, we do not use any action $a_{\theta j}$ more than $\texttt{OPT}_{\theta j}$ times. Observe that when $\mathcal{S}_t$ does not hold true

- the probability that we take an action $a_{\theta j}$ where $\texttt{DLP}_{\theta j} = 0$ is equal to 0, and
- the probability that we take an action $a_{\theta j}$ where $\texttt{DLP}_{\theta j} > 0$ is at most $\texttt{DLP}_{\theta j} + \ell \frac{[\log(t)]^{\frac{1+\epsilon}{2}}}{t^{\frac{1-\epsilon}{2}}}$.

The former is thus guaranteed to not be used more than $\texttt{OPT}_{\theta j}$ times; the latter is stochastically dominated by $\sum_{t=1}^{\hat{T}} Ber\left(\texttt{DLP}_{\theta j} + \ell \frac{[\log(t)]^{\frac{1+\epsilon}{2}}}{t^{\frac{1-\epsilon}{2}}}\right)$. We want to show that this is less than $\texttt{OPT}_{\theta j}$, with sufficiently high probability, to prove the lemma. Recall first (Lemma 3) that $\texttt{OPT}_{\theta j} \ge T\times\texttt{DLP}_{\theta j} - [T\log(T)]^{\frac{1+\epsilon}{2}}$ with probability at least $1-2kT^{\frac{-T^\epsilon}{\delta^2}}$. On the other hand, we have

$$\mu_{\theta j} := \mathbb{E}\left[\sum_{t=1}^{\hat{T}} Ber(\texttt{DLP}_{\theta j} + \ell \frac{[\log(t)]^{\frac{1+\epsilon}{2}}}{t^{\frac{1-\epsilon}{2}}})\right] \le \hat{T}\times\texttt{DLP}_{\theta j} + \ell\log(T)\int_1^{\hat{T}} \frac{1}{t^{\frac{1-\epsilon}{2}}} dt = \hat{T}\texttt{DLP}_{\theta j} + \ell\log(T)\frac{2\hat{T}^{\frac{1+\epsilon}{2}}}{\epsilon+1}.$$

Now, we can bound for an action $a_{\theta j}$ with $\mathtt{DLP}_{\theta j} > 0$

$$\begin{aligned}&\mathbb{P}\left[\sum_{t=1}^{\hat{T}} Ber(\mathtt{DLP}_{\theta j} + \ell\frac{[\log(t)]^{\frac{1+\epsilon}{2}}}{t^{\frac{1-\epsilon}{2}}}) > \mathtt{OPT}_{\theta j}\right] \\ &\leq 2kT^{\frac{-T^{\epsilon}}{\delta^2}} + \mathbb{P}\left[\sum_{t=1}^{\hat{T}} Ber(\mathtt{DLP}_{\theta j} + \ell\frac{[\log(t)]^{\frac{1+\epsilon}{2}}}{t^{\frac{1-\epsilon}{2}}}) > T\times\mathtt{DLP}_{\theta j} - [T\log(T)]^{\frac{1+\epsilon}{2}}\right].\end{aligned}$$

Suppose next that

$$T\times\mathtt{DLP}_{\theta j} - [T\log(T)]^{\frac{1+\epsilon}{2}} > \hat{T}\times\mathtt{DLP}_{\theta j} + \ell\log(T)\frac{2\hat{T}^{\frac{1+\epsilon}{2}}}{\epsilon+1} + [T\log(T)]^{\frac{1+\epsilon}{2}}$$

$$\text{or equivalently } (T-\hat{T})\times\mathtt{DLP}_{\theta j} \geq \ell\log(T)\frac{2\hat{T}^{\frac{1+\epsilon}{2}}}{\epsilon+1} + 2\left[T\log(T)\right]^{\frac{1+\epsilon}{2}}.$$

Observe that the left-hand side grows as $\Omega(T^{\frac{1}{2}+\epsilon})$, whereas the right-hand side grows as $\tilde{O}(T^{\frac{1+\epsilon}{2}})$, so for large enough T this holds true. We set T_0 to be the smallest constant such that this holds for $T \geq T_0$ for every θ and j with $\mathtt{DLP}_{\theta j} > 0$. Then, we can bound the above probability as

$$\begin{aligned}&2kT^{\frac{-T^{\epsilon}}{\delta^2}} + \mathbb{P}\left[\sum_{t=1}^{\hat{T}} Ber(\mathtt{DLP}_{\theta j} + \ell\frac{[\log(t)]^{\frac{1+\epsilon}{2}}}{t^{\frac{1-\epsilon}{2}}}) > \hat{T}\times\mathtt{DLP}_{\theta j} + \ell\log(T)\frac{2\hat{T}^{\frac{1+\epsilon}{2}}}{\epsilon+1} + [T\log(T)]^{\frac{1+\epsilon}{2}}\right] \\ &\leq 2kT^{\frac{-T^{\epsilon}}{\delta^2}} + \mathbb{P}\left[\sum_{t=1}^{\hat{T}} Ber(\mathtt{DLP}_{\theta j} + \ell\frac{[\log(t)]^{\frac{1+\epsilon}{2}}}{t^{\frac{1-\epsilon}{2}}}) - \mu_{ij} > [T\log(T)]^{\frac{1+\epsilon}{2}}\right] \\ &\leq 2kT^{\frac{-T^{\epsilon}}{\delta^2}} + e^{-2\frac{(T\log(T))^{1+\epsilon}}{\hat{T}}} \leq 2kT^{\frac{-T^{\epsilon}}{\delta^2}} + T^{-2},\end{aligned}$$

where the last line follows from Hoeffding's Inequality (see Proof of Lemma 1). Taking a union bound over all $a_{\theta j}$ with $\mathtt{DLP}_{\theta j} > 0$ — noting that we only need the bound from Lemma 3 once — we obtain a bound of $2kT^{\frac{-T^{\epsilon}}{\delta^2}} + k\ell T^{-2} < \frac{1}{T}$ for $T > \max\{2\delta^{\frac{2}{\epsilon}}, 2k\ell, T_0\}$. The result follows. □

Proof of Lemma 6. We argue similarly to the proof of Lemma 3: by Proposition 1, in period t,

$$\max_{\theta,j} |\mathtt{OPT}^t_{\theta j} - \mathtt{TDLP}_{\theta j}(t)| \leq \delta\max_{\theta} |\frac{\bar{t}}{t}N_\theta(t) - \bar{N}_\theta(t)|,$$

$$\text{implying } \forall\theta, j: \mathtt{OPT}^t_{\theta j} \geq \mathtt{TDLP}_{\theta j}(t) - \delta\max_{\theta} |\frac{\bar{t}}{t}N_\theta(t) - \bar{N}_\theta(t)|.$$

Then, by triangle inequality, $\delta\max_\theta |\frac{\bar{t}}{t}N_\theta(t) - \bar{N}_\theta(t)| \leq \delta\left(\max_\theta |\frac{\bar{t}}{t}N_\theta(t) - p_\theta\bar{t}| + \max_\theta |p_\theta\bar{t} - \bar{N}_\theta(t)|\right)$, and we find, with $\alpha = \frac{\sqrt{\bar{t}}}{\delta^2}$ in Lemma 1, that the probability of either term being greater $\frac{\bar{t}^{\frac{3}{4}}\sqrt{\log(\bar{t})}}{\delta}$ is at most $4k\bar{t}^{-2\sqrt{\bar{t}}/\delta^2}$, where we are using that $t > \bar{t}$ which follows from $\hat{T} > T/2$. This implies the result, i.e., we have with probability at least $1 - 4k\bar{t}^{-2\sqrt{\bar{t}}/\delta^2}$ that $\forall\theta, j: \mathtt{OPT}^t_{\theta j} \geq \mathtt{TDLP}_{\theta j}(t) - 2\bar{t}^{\frac{3}{4}}\sqrt{\log(\bar{t})}$. □

A.3. Proof of Lemma 7

As stated before, the lemma follows from the Berry-Esseen Theorem (Berry 1941), of which we use the following formulation.

PROPOSITION 6 **(Corollary 3 in Shevtsova (2011)).** *Let $X_1, \ldots, X_\tau$ be i.i.d. random variables drawn from a distribution fulfilling $\mathbb{E}[X_i] = 0$, $\sigma^2 = Var[X_i] = E[X_i^2] = 1$, and finite third absolute moment β_3. Let $S_\tau = \sum_i X_i$, $Z \in N(0,1)$, and $\phi(r)$ be the CDF of Z. Then*

$$\Delta_\tau = \sup_r \left|\phi(r) - \mathbb{P}\left[S_\tau < r\sqrt{\tau}\right]\right| \leq 0.4748\beta_3/\sqrt{\tau}$$

holds for all n.

Define the Bernoulli random variables $X_1, \ldots, X_\tau$ with

$$X_i = \begin{cases} \sqrt{\frac{1-p_\theta}{p_\theta}} & \text{with probability } p_\theta \\ -\sqrt{\frac{p_\theta}{1-p_\theta}} & \text{with probability } 1-p_\theta \end{cases}$$

Observe that $\mathbb{E}[X_i] = 0$, $\sigma^2 = 1 - p_\theta + p_\theta = 1$, and $\beta_3 = \mathbb{E}[|X^3|] \leq \frac{1}{\sqrt{p_\theta}} + \frac{1}{\sqrt{1-p_\theta}}$. Then, with S_n, Z as in Proposition 6 we have

$$\sup_r |\phi(r) - \mathbb{P}[S_\tau < r\sqrt{\tau}]| \leq \left(\frac{1}{\sqrt{p_\theta}} + \frac{1}{\sqrt{1-p_\theta}}\right)\frac{0.4748}{\sqrt{\tau}}.$$

For fixed $r > 0$, we have $1 > \phi(r) > 0$, so for large enough τ there exists $\bar{\xi}_r \in (0,1)$ with

$$\bar{\xi}_r := \phi(r) - \left(\frac{1}{\sqrt{p_\theta}} + \frac{1}{\sqrt{1-p_\theta}}\right)\frac{0.4748}{\sqrt{\tau}} \leq \mathbb{P}[S_\tau < r\sqrt{\tau}].$$

Setting $\xi_r = 1 - \bar{\xi}_r$ we find that $\mathbb{P}[S_\tau \geq r\sqrt{\tau}] \geq \xi_r > 0$. Now, define $Y_i = X_i\sqrt{p_\theta(1-p_\theta)} + p_\theta$, to obtain a random variable that is 1 with probability p_θ and 0 otherwise. Further, we have

$$N_\theta \stackrel{d}{=} \sum_i Y_i = \sum_i X_i\sqrt{p_\theta(1-p_\theta)} + p_\theta = np_\theta + \sqrt{p_\theta(1-p_\theta)}S_\tau$$

With probability at least $\frac{\xi_r}{\sqrt{p_\theta(1-p_\theta)}}$ we have $S_\tau \geq r\sqrt{\tau}$, and thus, with probability at least $\frac{\xi_m}{\sqrt{p_\theta(1-p_\theta)}}$, we have $N_\theta \geq np_\theta + \sqrt{p_\theta(1-p_\theta)}$. □

A.4. Proofs of further impossibility results

Proof of Proposition 3. Consider a bin packing instance as in Proposition 2 in which bins have size 3 and items are of size either 1 or 2 with probability $\frac{1}{2}$; the time horizon T is a geometric random variable with mean τ. As in Proposition 2, all arriving items of size 2 need to be put into a bin of configuration $\{1, 2\}$

whereas items of size 1 can either be in a $\{1,2\}$ configuration or in a $\{1,1,1\}$ configuration. The main observation necessary to derive the result is that the optimal Bellman equation gives a threshold solution that sets a threshold d_τ and puts arriving items of size 1 into bins of configuration $\{1,1,1\}$ when there are d_τ bins with only one 1 in them (such bins can have an arriving 2 added to them), i.e., it acts as follows:

- when an item of size 2 arrives, and there is a bin with only one 1 in it, add the size 2 item to that bin; else, add the arriving size 2 item to a new bin;
- when an item of size 1 arrives, and there is a bin with only one 2 in it, add the arriving size 1 item to that bin; if there is no such bin, but there is a bin that has two size 1 items in it, add the arriving size 1 item to that one; if there is no such bin either, and there are less than d_τ bins that have just a single item of size 1 in them, add the arriving size 1 item to a new bin; else, add the arriving size 1 item to an existing bin that has a single 1 in it.

We distinguish between the case that $d_\tau \le \sqrt{\tau}$ and $d_\tau > \sqrt{\tau}$. Further, we condition on the time horizon lasting between τ and 2τ periods, which occurs with constant probability (as τ grows large). Then, the following events each occur with constant probability:

1. over the entire horizon the number of items of size 1 is $2\sqrt{\tau}$ greater than the number of items of size 2;
2. over the first $\tau/2$ periods there are at least $2\sqrt{\tau}$ more arrivals of size 1 than of size 2 but over the entire horizon there are more arrivals of size 2 than of size 1.

With $d_\tau \le \sqrt{\tau}$, under the first event, at the end of the horizon there are $\Omega(\sqrt{\tau})$ items of size 1 in bins by themselves; these could be in a third as many bins of configuration $\{1,1,1\}$; with $d_\tau > \sqrt{\tau}$, under the second event, the algorithm creates $\Omega(\sqrt{\tau})$ bins of configuration $\{1,1,1\}$ in the first $\tau/2$ periods, whereas the optimal solution requires none of these; thus, for any threshold d_τ the expected loss is $\Omega(\sqrt{\tau})$. $\square$

Proof of Proposition 4. We consider the following events:

- E_1 is the event that there are at least $\sqrt{T}/2$ items of value 2 in both the first and the last $T/2$ periods;
- E_2 is the event that there are at most $T/4$ items of value 3 over the first $T/2$ periods;
- E_3 is the event that there are at least $T/2$ items of value 3 over the entire horizon;
- E_4 is the event that there are at most $T/2 - \sqrt{T}$ items of value 3 over the entire horizon.

Observe that both $E_1 \cap E_2 \cap E_3$ and $E_1 \cap E_2 \cap E_4$ occur with constant probability by Lemma 7; now, consider the number of items of value 2 accepted by an algorithm in the first $T/2$ periods: if it is at least $\sqrt{T}/4$, then the algorithm incurs $\Omega(\sqrt{T})$ loss under $E_1 \cap E_2 \cap E_3$, and if it is at most $\sqrt{T}/4$ then the algorithm incurs $\Omega(\sqrt{T})$ loss under $E_1 \cap E_2 \cap E_4$. Thus, in this instance any algorithm must incur $\Omega(\sqrt{T})$ loss. □

Appendix B: Application to online packing problems

In this section we describe how our framework captures online packing problems or variants thereof, including AdWords (Alaei et al. 2013), Network Revenue Management (Talluri and Van Ryzin 1998), or refugee resettlement (Bansak and Paulson 2024). In these problems, for a fixed time horizon T, we have a budget of resources $\vec{B}$. Upon each arrival of type θ we need to make an irrevocable assignment decision, wherein an arrival is assigned to a product j, consumes some resources $A_{\theta j}$ for each resource type j, and yields some value $c^1_{\theta j}$ — note that the latter two are dependent on both the arrival type and the product assignment; to fit the minimization objective description in Section 2 we will minimize $-\vec{c}^1 \cdot \vec{x}$, which maximizes the value. In the canonical quantity-based network revenue management problem, the arrival is a specific request and the decision is just *accept/reject*; we can model the *reject* decision by allowing for a product with infinite resources and no value. In refugee resettlement there is no *reject* decision as arriving refugees need to be send to a location within a host country; the value to the optimization problem comes from the probability they find employment, as captured through an exogenous machine learning model. As a relaxation of these problems we allow $f_{\vec{N}}$ to violate the budget constraint at a penalty which is greater than the largest value; when T is known from the beginning our algorithm would never violate the budget/incur that penalty; when T is a priori unknown our algorithm may, with vanishingly small probability under Assumption 3, violate the budget and incur the penalty. We have z_j as a decision variable capture by how much the resource constraint on resource j has been violated, and have a penalty of $c^2_j = \max\{c_{\theta j} + 1\}$ Then, we can write the objective as follows

$$
\begin{aligned}
f_{\vec{N}}(\vec{x}) \quad = \quad & \min_{\vec{z}} -\vec{c}^{\,1} \cdot \vec{x} + \vec{c}_2 \cdot \vec{z} \\
\text{s.t.} \quad & \sum_j x_{\theta j} = N_\theta \, \forall \theta \\
& A\vec{x} - I\vec{z} \le \vec{B} \sum_\theta N_\theta \\
& z_j \ge 0 \; \forall j,
\end{aligned}
$$

Observe that for any fixed $\vec{x}$ the first term in the objective captures the *negative of the value of the product assignments* (which we aim to minimize), whereas the optimal $z_j = \max\{0, \sum_\theta A_{\theta j} x_{\theta j} - B_j\}$ captures by how much the resource constraint on resource j is violated, and penalizes it by a number greater than the value of any resource assignment.

Egalitarian welfare

A particular case of interest is the maximization of egalitarian welfare discussed by Balseiro and Xia (2022). In their model, each offline agent $i \in \{1,\dots,n\}$ obtains utility u^i_θ when they receive an item of type θ; thus, with $x_{\theta i}$ items of type θ assigned to agent i, the utility of agent i is $\sum_\theta x_{\theta i} u^i_\theta$. The goal is to maximize the minimum utility, i.e., to maximize $\min_i \sum_\theta x_{\theta i u^i_\theta}$ subject to $\sum_i x_{\theta i} = N_\theta$. To translate this into our formulation, we first observe that the minimum utility, for a given $\vec{x}$ with $\sum_i x_{\theta i} = N_\theta$, can be written as

$$\min_{\vec{z}} -z_0 \qquad \text{s.t.} \qquad z_i - \sum_\theta x_{\theta i} u^i_\theta \le 0 \quad \forall i \in \{1,\dots,n\} \qquad z \ge 0, \qquad z_0 - z_i \le 0 \quad \forall i \in \{1,\dots,n\}.$$

The first constraint encodes that z_i is bounded above by the utility of agent i; the last constraint ensures that z_0, which the objective maximizes, is no larger than any agent's utilities. In other words, z_0 is the minimum utility among the agents. To express this through the function $f_{\vec{N}}(\vec{x})$ described in Section 2, we set the following parameters:

- $\vec{c}_1 = \vec{0}$ and $\vec{B} = 0$;
- $\vec{c}_2$ is a $(n+1)$-dimensional vector with a -1 as its first entry and all other entries equal to 0;
- A is such that the ith row of $A\vec{x}$ equals $-\sum_\theta x_{\theta i} u^i_\theta$ for $i = 1,\dots,n$, whereas all other rows are 0;
- D is such that the ith row of $D\vec{z}$ is z_i for $i = 1,\dots,n$ and it is $z_0 - z_{i-n}$ for $i = n+1,\dots,2n$.

Appendix C: Known horizon and lazy resolving

In this appendix we consider different assumptions on the arrival process, the information structure on the length of the time-horizon, and the objective function. Specifically, we focus on a setting where the objective function (not necessarily linear) fulfills properties (i) and (ii) in Proposition 1, the time horizon T is known from the beginning, we have one sample path (a *trace*) of T arrivals, and the arrivals come from a distribution that fulfills the following two assumptions.

Assumption 4 **(Anytime Concentration)**. *For each type $\theta \in \Theta$ and any period t, we assume that the number of future arrivals $\bar{N}_\theta(t)$ is concentrated about its mean; formally, $\exists$ constant $\alpha > 0$, independent of t, such that $\forall\, \theta, t$ and every $x > 0$*

$$\mathbb{P}\left[\left|\bar{N}_\theta(t) - \mathbb{E}[\bar{N}_\theta(t)]\right| > x\right] \le 2e^{\frac{-\alpha x^2}{t}}.$$

Assumption 5 **(Anytime Thickness)**. *For each type $\theta \in \Theta$ and any period t, we assume that the expected number of future arrivals $\bar{N}_\theta(t)$ grows faster than $\sqrt{t}$; formally, $\exists$ positive constants β, γ independent of t such that $\forall \theta, t$ and every $x > 0$*

$$\mathbb{E}\left[\bar{N}_\theta(t)\right] \geq \beta \bar{t}^{0.5+\gamma}.$$

Observe that Assumption 1, with $\gamma = 1/2$, is a special case of Assumption 5. Proposition 7 shows that Assumption 5 is tight in the sense that one cannot guarantee an expected loss of $o(\sqrt{T})$ when $\gamma = 0$. Also, note that Lemma 1 (implied by Assumption 1) is a variant of Assumption 4.

Algorithm 3:

– Initialize $x^{\texttt{ALG}}_{\theta j}(0) = 0; F_{\theta j} = 0 \ \forall \theta, j$

for $t = 1, \ldots, T$ **do**

  if $F_{\theta_t j} < 1$ *for every action* j **then**

    – Solve optimization problem $\texttt{SDLP}(t)$;

    – Set $F_{\theta j} = \left\lfloor \texttt{SDLP}_{\theta j}(t) - x^{\texttt{ALG}}_{\theta j}(t) - 2\bar{t}^{\left(\frac{1+\gamma}{2}\right)} \sqrt{\log \bar{t}} \right\rfloor \ \forall \theta, j$

  end

  – In period t, use action $a_{\theta_t j}$ for some $j \in \arg\max_j \{F_{\theta_t j}\}$; reduce $F_{\theta_t j}$ by 1;

  – Set $x^{\texttt{ALG}}(t)$ by incrementing $x^{\texttt{ALG}}_{\theta_t j}(t-1)$ by 1 and leaving all other values unchanged;

end

We denote the sample path we have access to by $\hat{\theta}_1, \ldots, \hat{\theta}_T$, write $\widetilde{N}^f_\theta(t)$ for the number of type θ arrivals in periods $t+1, \ldots, T$ on that sample path, and $\widetilde{N}^f$ for the vector of those arrivals over the entire sample path; it is worth noting that having such a sample path is a significantly weaker assumption than full distributional information, as the later always allows one to draw one sample. Relative to the i.i.d. setting in the main body, this setting here allows for limited correlations between arrivals in different periods. Similar to Algorithm 1 solving TDLP once the time horizon is known, Algorithm 3 relies on solving the following:

$$\min_{\vec{x}} f_{\widetilde{N}^f}(\vec{x}) \qquad \text{s.t. } \geq \vec{x}^{\texttt{ALG}}(t-1). \tag{SDLP(t)}$$

Again, similar to Algorithm 1, Algorithm 3 has the following property.

THEOREM 2. *Suppose the arrivals fulfill Assumption 4 and 5 and the objective fulfills properties (i) and (ii) of Proposition 1; then there exists an algorithm* **ALG** *and a constant $M_2(\beta,\lambda,\delta,\epsilon,\ell)$, independent of T, such that the loss of Algorithm 3 can be bounded as $\mathcal{L}_{\texttt{ALG}} \le M_2(\beta,\delta,\lambda,\epsilon,\ell)$.*

Observe that Theorem 2 does not rely on Assumption 2. Thus, we do not require a solution $\texttt{OPT}^t$ to be unique below.[12] The proof of Theorem 2 requires us to show the following result, similar to Lemma 6.

LEMMA 8. *Consider a period t with actions $x^{ALG}_{\theta j}(t-1)$ already taken in periods $[1,t-1]$; with probability at least $1-4\bar{t}^{-\frac{\alpha \bar{t}^{\gamma}}{\delta^2}}$ there exists some solution OPT^t with $OPT^t_{\theta j} \ge \left\lfloor SDLP_{\theta j}(t) - 2\bar{t}^{\left(\frac{1+\gamma}{2}\right)}\sqrt{\log \bar{t}}\right\rfloor \forall \theta, j$.*

Proof sketch of Theorem 2. Denote, as in the proof of Theorem 1, by $T_1 = \hat{T}, T_2, \ldots, T_s$ the periods in which Algorithm 3 solves $\texttt{SDLP}(t)$, by T_η the first period in which Algorithm 3 solves $\texttt{TDLP}(t)$ and it is the case that for some θ all budgets $F_{\theta j} < 1$, and let $\mathcal{T}_i$ be the event that the LCB created in period T_i fails to hold true for any optimal solution $\texttt{OPT}^{T_i}$:

$$\mathcal{T}_i = \left\{ \forall\, \texttt{OPT}^{T_i}\ \exists \theta, j : \texttt{OPT}^{T_i}_{\theta j} < \left\lfloor \texttt{TDLP}_{\theta j}(T_i) - 2\bar{T}_i^{\frac{3}{4}}\sqrt{\log \bar{T}_i}\right\rfloor \right\}.$$

Then we can bound

$$\mathcal{L}_{\texttt{ALG}} = \mathbb{E}\left[\texttt{OPT}^T - \texttt{OPT}^1\right] \le 2\lambda \mathbb{E}\left[\sum_{i=1}^{\eta-1}(T_{i+1}-T_i)\mathbb{P}[\mathcal{T}_i] + \bar{T}_\eta\right] \le 2\lambda \mathbb{E}\left[\sum_{i=1}^{\eta-1}\bar{T}_i\mathbb{P}[\mathcal{T}_i] + \bar{T}_\eta\right].$$

We bound the sum similarly to how we did in the proof of Theorem 1 (using Lemma 8 instead of Lemma 6, where with $\gamma = 1/2$ it is the same bound).

For $\mathbb{E}\left[\bar{T}_\eta\right]$ we construct an upper bound by noting that in period t, $\texttt{SDLP}$ assigns $\widetilde{N}^f_\theta(t)$ actions for type θ, meaning at least one action is assigned at least $\widetilde{N}^f_\theta(t)/\ell$ actions; with high probability (by Assumption 4) this is at least $\frac{\mathbb{E}[\bar{N}_\theta(t)]}{2\ell} \ge \beta \bar{t}^{.5+\gamma}$ which, for $\bar{t}$ large enough is greater $2\bar{t}^{\left(\frac{1+\gamma}{2}\right)}\sqrt{\log \bar{t}}$. $\square$

Proof of Lemma 8. Consider the solution $\texttt{OPT}^t$ to $\texttt{SDLP}(t)$ that minimizes $\max_{\theta,j}|\texttt{OPT}^t_{\theta j} - \texttt{SDLP}_{\theta j}(t)|$, where $\texttt{SDLP}(t)$ is the solution Algorithm 3 finds when solving in period t. Then,

$$\mathbb{P}\left[\exists \theta, j : \texttt{OPT}^t_{\theta j} < \lfloor \texttt{SDLP}_{\theta j}(t) - 2\bar{t}^{\left(\frac{1+\gamma}{2}\right)}\sqrt{\log \bar{t}}\rfloor\right] \le$$
$$\mathbb{P}\left[\max_{\theta,j}|\texttt{OPT}^t_{\theta j} - \texttt{SDLP}_{\theta j}(t)| \ge 2\bar{t}^{\left(\frac{1+\gamma}{2}\right)}\sqrt{\log \bar{t}}\right] \le$$

[12] Recall that for functions $f_{\vec{N}}$ as described in Section 2 this was without loss of generality, but for the more general class of functions we discuss here this need not hold true.

$$\mathbb{P}\left[\max_{\theta}|\bar{N}_{\theta}(t)-\widetilde{N}^{f}_{\theta}(t)|\geq\frac{2\bar{t}^{\left(\frac{1+\gamma}{2}\right)}\sqrt{\log\bar{t}}}{\delta}\right]\leq$$

$$k\max_{\theta}\mathbb{P}\left[|\bar{N}_{\theta}(t)-\widetilde{N}^{f}_{\theta}(t)|\geq\frac{2\bar{t}^{\left(\frac{1+\gamma}{2}\right)}\sqrt{\log\bar{t}}}{\delta}\right]\leq$$

$$k\max_{\theta}\left(\mathbb{P}\left[|\bar{N}_{\theta}(t)-\mathbb{E}[\bar{N}_{\theta}(t)]|\geq\frac{\bar{t}^{\left(\frac{1+\gamma}{2}\right)}\sqrt{\log\bar{t}}}{\delta}\right]+\mathbb{P}\left[|\mathbb{E}[\bar{N}_{\theta}(t)]-\widetilde{N}^{f}_{\theta}(t)|\geq\frac{\bar{t}^{\left(\frac{1+\gamma}{2}\right)}\sqrt{\log\bar{t}}}{\delta}\right]\right)\leq$$

$$4e^{-\frac{\alpha\bar{t}^{\gamma}\log(\bar{t})}{\delta^{2}}}=4\bar{t}^{-\frac{\alpha\bar{t}^{\gamma}}{\delta^{2}}}\quad\square$$

Lazy resolving. In this part we argue that Algorithm 3 has not only strong performance guarantees, but also allows for significant lesser resolving; in particular, we claim that with i.i.d. arrivals it only needs to resolve `SDLP`, in expectation, $O(\log\log(T))$ many times. To see this, note that after resolving in period t, for large enough $\bar{t}$, we have $\bar{t}-o(\bar{t})$ periods to go before we next need to resolve. Indeed, it next resolves after observing, for some type θ, at least $\widetilde{N}^{f}_{\theta}(t)-\ell 2\bar{t}^{\frac{3}{4}}\sqrt{\log\bar{t}}$ arrivals. Now, with i.i.d. arrivals we find that $\widetilde{N}^{f}_{\theta}(t)$ is close to, or greater, $\bar{t}p_{\theta}$ with high probability (in $\bar{t}$, by Assumption 4 or Lemma 1); thus, θ requires the algorithm to resolve next after observing about $\bar{t}p_{\theta}-o(\bar{t})$ more arrivals of type θ; however, before observing $\bar{t}p_{\theta}-o(\bar{t})$ arrivals of type θ, with high probability, we observe $\bar{t}p_{\theta'}-o(\bar{t})$ of each other type θ'; thus, only $o(\bar{t})$ periods remain when the algorithm next resolves.

Necessity of Assumption 5 In this section we prove that Assumption 5 is tight in the sense that with $\gamma=0$ one may not be able to obtain the general uniform loss guarantees of Theorem 2.

PROPOSITION 7. *There exists a family of instances wherein $\mathbb{E}[\bar{N}_{\theta}(t)]\geq\beta\bar{t}^{0.5}$ holds true $\forall t$, yet no online algorithm can obtain uniform loss to the online knapsack problem.*

Proof. Consider a time horizon of even length T and a budget $T/2$. Items have unit size and value 1, 2, or 3. For the first $T/2$ periods, all arriving items have value 2. Then, there are $T/2-\sqrt{T}$ arrivals of value either 1 or 3 with probability $\frac{1}{2}$ each. Finally, the last $\sqrt{T}$ arrivals have value 1, 2, or 3 with probability $\frac{1}{3}$ each. Observe that with this family of instances we have $\mathbb{E}[\bar{N}_{\theta}(t)]\geq\frac{\sqrt{\bar{t}}}{3}$ since $\mathbb{E}[\bar{N}_{\theta}(t)]=\frac{\bar{t}}{3}$ for $t\geq T-\sqrt{T}$, and $\mathbb{E}[\bar{N}_{\theta}(t)]\geq\frac{\sqrt{T}}{3}\geq\frac{\sqrt{\bar{t}}}{3}$ for $t<T-\sqrt{T}$, so these instances fulfill the assumption of the proposition with $\beta=1/3$. Further, by Lemma 7 there exists $\xi>0$ such that

$$\mathbb{P}[\bar{N}_{2}(T/2)+\bar{N}_{3}(T/2)]+\leq\frac{T}{4}-\sqrt{T}]>\xi \text{ and } \mathbb{P}[\bar{N}_{3}(T/2)\geq\frac{T}{4}+\sqrt{T}]>\xi.$$

Now, consider an algorithm `ALG` that accepts $x^{\mathtt{ALG}}_{2}(\frac{T}{2})$ items of value 2 in the first $\frac{T}{2}$ periods. We distinguish between two cases.

Case $x_2^{ALG}(\frac{T}{2}) \geq \frac{T}{4}$. With probability at least ξ we have $\bar{N}_3(T/2) \geq \frac{T}{4} + \sqrt{T}$, i.e., there are at least $\frac{T}{4} + \sqrt{T}$ arrivals of value 3 in the second half of the time horizon, of which `ALG` can only accept $\frac{T}{4}$ yielding an expected loss, relative to the clairvoyant, of at least $\xi\sqrt{T} \in \Omega(\sqrt{T})$.

Case $x_2^{ALG}(\frac{T}{2}) < \frac{T}{4}$ With probability at least ξ we have $\bar{N}_2(T/2) + \bar{N}_3(T/2) + \leq \frac{T}{4} - \sqrt{T}$, i.e., there are at most $\frac{T}{4} - \sqrt{T}$ arrivals of value either 2 or 3 in the remaining periods. Thus, `ALG` must accept at least $\sqrt{T}$ items of value 1 or leave $\sqrt{T}$ of its budget unused, yielding an expected loss, relative to the clairvoyant, of at least $\xi\sqrt{T} \in \Omega(\sqrt{T})$. $\square$

Appendix D: Explicit δ-bounds

In this appendix we provide explicit analytic bounds on the Lipschitz constant δ from Proposition 1 for two specific examples in our paper.

Online payment routing problem from Section 5

We begin with the online payment routing problem in Section 5. Here, we have $v_\theta \in \{2, 5, 8\}$, Provider A charges $b_A = 0.40\$$, and $m_A = a_A = 0\$$ per payment, i.e., processing each payment costs 0.40$ regardless of the transaction value, whereas provider B charges a proportional $m_B = 0.10\$$ and $b_B = 0\$$ per dollar processed (e.g., processing a 8$ payment would cost 0.80$ with that provider) while also requiring a minimum average amount of 0.30$ per payment. For given N_2, N_5, N_8 we can write the objective of a solution explicitly as follows:

$$\begin{aligned} & \sum_\theta .4x_{\theta A} + \max\left\{.3\sum_\theta x_{\theta B}, .1\sum_\theta \theta x_{\theta B}\right\} \\ \text{(since } x_{\theta A} + x_{\theta B} = N_\theta\text{)} \qquad &= \sum_\theta .4x_{\theta A} + \max\left\{.3\sum_\theta (N_\theta - x_{\theta A}), .1\sum_\theta \theta(N_\theta - x_{\theta A})\right\} \\ &= \sum_\theta .1x_{\theta A} + \max\left\{.3\sum_\theta N_\theta, \sum_\theta \Big(.1\theta(N_\theta - x_{\theta A}) + .3x_{\theta A}\Big)\right\} \end{aligned}$$

We first note that with $x_{\theta A} = N_\theta$ for every θ, our objective becomes $.4\sum_\theta N_\theta$. In that case, the maximum evaluates to $.3\sum_\theta N_\theta$ with the first and second term in the maximum being equal at the lowest value the maximum can attain. Now, in order to lower the value of a solution, we need to lower some of the $x_{\theta A}$. Suppose we decrease each $x_{\theta A}$ by Δ_θ. Then the first term decreases by by $.1\sum_\theta \Delta_\theta$ and the max increases by $\max\{0, \sum_\theta \Delta_\theta(.1\theta - .3)\}$. In other words, we must be in one of the following cases, depending on whether the maximum evaluates to more or less than $.3\sum_\theta N_\theta$.

More than $.3\sum_\theta N_\theta$*.* If the maximum evaluates to more than $.3\sum_\theta N_\theta$ then the derivative of the objective with respect to $x_{\theta A}$ is $-.1\theta + .4$, so we have $x_{\theta A} = 0$ for $\theta < 4$ and $x_{\theta A} = N_\theta$ for $\theta > 4$. It follows that the objective evaluates to $.3N_2 + .5N_5 + .8N_8$.

At most $.3\sum_\theta N_\theta$*.* In this case $\sum_\theta \Delta_\theta$ must be maximized over all values of Δ_θ such that $\sum_\theta \Delta_\theta(.1\theta - .3) \leq 0$. Clearly, in order to find the best way to set such Δ we can first increase Δ_θ for all θ such that $.1\theta - .3 \leq 0$ and then continue by increasing the smallest θ such that $.1\theta - .3 > 0$ until either $\sum_\theta \Delta_\theta(.1\theta - .3) = 0$ or $\Delta_\theta = N_\theta \forall\theta$. In particular, for our example we obtain

$$x_{2A} = 0 \qquad x_{5A} = \min\{.5N_2, N_5\} \qquad x_{8A} = \min\{\max\{(N_2 - 2N_5)/5, 0\}, N_8\}.$$

As this gives a strictly better objective than the above case (whenever $N_2 \geq 2, N_5 \geq 1$), we can use the right-hand sides to obtain a bound on the Lipschitz constant δ for Proposition 1.

Though a larger number of types makes it much harder to write out the exact solution, we note that most of the reasoning here extends to the data-driven example in Section 5.

Bin packing instance from Example 1

We again aim to provide a closed-form solution for the problem with known N_θ. However, in contrast to the payment routing problem, the optimal solution to the bin packing instance is more intractable in general and we thus focus on arrival patterns as they would occur with high probability. Moreover, to simplify exposition, we will make some approximations where, e.g., just a single bin would be affected.

Recall that we have bins of size 10 and items of size 1,3,4,5, and 8 arrive in each period, with respective probabilities $p_1 = .25, p_3 = .25, p_4 = .125, p_5 = .25, p_8 = .125$. If T items arrived exactly as in expectation we could pack them in $T/8$ bins of configuration $\{1,1,8\}$, $T/8$ bins of configuration $\{3,3,4\}$, and $T/8$ bins of configuration $\{5,5\}$. We focus on $N_1, N_2, N_3, N_4, N_5, N_8$ that are approximately as in expectation. Also, to simplify notation, we will assume that all fractions evaluate as integer (noting that this only affects a constant number of bins which would not affect our guarantees).

Case $N_8 \geq N_1/2$**.** When the number of items of size 8 is at least half the number of items of size 1, the optimal solution becomes very simple: all items of size 1 are in bins of configuration $\{1,1,8\}$; items of size 5 are in $N_5/2$ bins of configuration $\{5,5\}$; finally, we have $\min\{N_3/2, N_4\}$ bins of configuration $\{3,3,4\}$ and depending on whether $N_4 < N_3/2$ or $N_4 > N_3/2$ we have either $(N_4 - N_3/2)/2$ bins of configuration $\{4,4\}$ or $(N_3 - 2\times N_4)/3$ bins of configuration $\{3,3,3\}$.

Case $N_8 < N_1/2$. In this case, we have $\hat{N}_1 := N_1 - 2N_8$ items of size 1 that do not fit with items of size 8 in bins of configuration $\{1,1,8\}$. Depending on whether $N_4 < N_3/2$ or $N_4 > N_3/2$, these are optimally placed either with items of size 3 or with items of size 4:

- If $N_4 < N_3/2$, then we have $\hat{N}_3 := (N_3 - 2 \times N_4)$ items of size 3 that are not in bins of configuration $\{3,3,4\}$ and $\hat{N}_1$ items of size 1 that are not in bins of configuration $\{1,1,8\}$. These items can then be optimally placed in bins of configurations $J_1 = \{1,3,3,3\}, J_2 = \{1,\ldots,1\}$, where the number of bins of each of these two configuration is $x_1 = \hat{N}_3/3$ and $x_2 = \max\{\hat{N}_1 - x_1, 0\}/10$.
- If $N_4 > N_3/2$, then we have $\hat{N}_4 := (N_4 - N_3/2)$ items of size 4 that are not in bins of configuration $\{3,3,4\}$ and $\hat{N}_1$ items of size 1 that are not in bins of configuration $\{1,1,8\}$. These items are then placed in bins of configurations $J_1 = \{1,1,4,4\}, J_2 = \{1,\ldots,1\}$, where the number of bins of each of these two configurations is $x_1 = \hat{N}_4/2$ and $x_2 = \max\{\hat{N}_1 - 2x_1, 0\}/10$.

Again, for either case, we can use the right-hand side formulas for the number of bins of each configuration to determine a bound on the Lipschitz constant δ in Proposition 1.