Working Paper 23-047

Group Fairness in Dynamic Refugee Assignment

Daniel Freund
Thodoris Lykouris
Elisabeth Paulson
Bradley Sturt
Wentao Weng

Group Fairness in Dynamic Refugee Assignment

Daniel Freund
Massachusetts Institute of Technology

Thodoris Lykouris
Massachusetts Institute of Technology

Elisabeth Paulson
Harvard Business School

Bradley Sturt
University of Illinois at Chicago

Wentao Weng
Massachusetts Institute of Technology

Working Paper 23-047

Funding for this research was provided in part by Harvard Business School.

Group fairness in dynamic refugee assignment

Daniel Freund* Thodoris Lykouris† Elisabeth Paulson‡ Bradley Sturt§
Wentao Weng¶

Abstract

Ensuring that refugees and asylum seekers thrive (*e.g.*, find employment) in their host countries is a profound humanitarian goal, and a primary driver of employment is the geographic location within a host country to which the refugee or asylum seeker is assigned. Recent research has proposed and implemented algorithms that assign refugees and asylum seekers to geographic locations in a manner that maximizes the average employment across all arriving refugees. While these algorithms can have substantial overall positive impact, using data from two industry collaborators we show that the impact of these algorithms can vary widely across key subgroups based on country of origin, age, or educational background. Thus motivated, we develop a simple and interpretable framework for incorporating *group fairness* into the dynamic refugee assignment problem. In particular, the framework can flexibly incorporate many existing and future definitions of group fairness from the literature (*e.g.*, minmax, randomized, and proportionally-optimized within-group). Equipped with our framework, we propose two bid-price algorithms that maximize overall employment while simultaneously yielding provable group fairness guarantees. Through extensive numerical experiments using various definitions of group fairness and real-world data from the U.S. and the Netherlands, we show that our algorithms can yield substantial improvements in group fairness compared to state-of-the-art algorithms with only small relative decreases ($\approx$ 1%-2%) in global performance.

1 Introduction

Over two million new refugees and asylum seekers are projected to require resettlement in 2023 to escape violence and persecution in their origin countries [UNH22], and ensuring the successful integration of these displaced people is a profound humanitarian and policy goal. Successful integration (typically measured by gainful employment after resettlement[1]) is crucial both to enhance the quality of life of refugees and asylum seekers and to invigorate the local economy of host countries. The challenge of facilitating the integration of refugees and asylum seekers into society is typically the purview of non-profit and governmental resettlement agencies such as the Swiss State Secretariat of Migration (SEM) in Switzerland, the Central Agency for the Reception of Asylum Seekers (COA) in the Netherlands, and the Lutheran Immigration and Refugee Service (LIRS) in the United States.

*Massachusetts Institute of Technology, dfreund@mit.edu
†Massachusetts Institute of Technology, lykouris@mit.edu
‡Harvard University and Stanford University, epaulson@stanford.edu
§University of Illinois at Chicago, bsturt@uic.edu
¶Massachusetts Institute of Technology, wweng@mit.edu

[1] For example, in the US, the Refugee Act of 1980 mandates annual audits of proxies for this metric.

A primary driver for successful integration is the geographic location in a host country to which a refugee or asylum seeker is assigned. Traditionally, case officers within resettlement agencies have made assignments based on local quotas and their own judgment. In the past few years, however, innovations in analytics have given rise to machine learning (ML) models that predict integration outcomes using personal characteristics [BFH$^+$18]. The advent of ML in this context opens up new possibilities for using analytics to improve resettlement decisions.

Equipped with ML models, the dynamic refugee assignment problem faced by resettlement agencies can be formulated as follows. Over the course of an extended period of time (*e.g.*, a fiscal year), the resettlement agency sequentially receives new cases $t = 1, \ldots, T$, where each case consists of a family or individual refugee/asylum seeker. The agency uses ML to compute an estimate of the probability w_{tj} that case t would successfully find employment in each geographic location j. Each location j has a capacity s_j on the number of refugees or asylum seekers that it is expected to resettle over that period. For some cases, the assignment is pre-determined by existing family ties or other constraints such as medical or educational considerations. For the *free* cases—those that can be resettled to any location—the resettlement agency makes an irrevocable decision of which location to assign the case. This decision must not violate the capacity constraints at any geographic location, and the objective of the resettlement agency in this problem is to maximize the average employment rate across all cases.

Several recent works have tackled the dynamic refugee assignment problem via optimization. The first papers to propose the use of optimization and ML in this context ([BFH$^+$18] and later [AAM$^+$21]) study a static variant in which all refugees arrive simultaneously and develop exact algorithms based on network flows and integer programming. Subsequent works [AGP$^+$21, BP22] formulate a dynamic version of the problem where refugees arrive sequentially from a probability distribution and design efficient heuristics to overcome computational intractability. These algorithms yield significant predicted gains (up to 50% increases) in average employment rate when compared to assignment decisions under current practice.

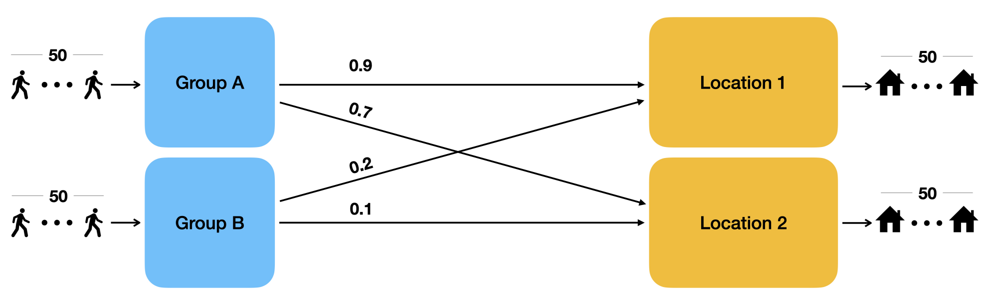

Figure 1: There are two groups $\{A, B\}$ with 50 refugees each and two locations $\{1, 2\}$ with capacities of 50. Each case of Group A finds employment with probability 0.9 when assigned to Location 1 and 0.7 when assigned to Location 2. For cases of Group B the probabilities are 0.2 and 0.1.

Although these algorithms represent a significant leap in the use of analytics to drive societal impact, one may worry that global optimization may consistently disadvantage a particular subgroup of refugees, *e.g.*, subgroups defined by country of origin, age, or educational background. For example, consider the setting of Figure 1, where members of groups A and B each achieve higher employment in Location 1. In this example, global optimization (i.e., optimizing for average employment across all cases) yields outcomes that are clearly unfair for any reasonable definition of group fairness, as all members of Group B are assigned to their least preferred location (Location 2) while all members of Group A are assigned to their most preferred location (Location 1).

Why might group fairness be of central importance in the setting of dynamic refugee assignment? First, the use of a refugee assignment algorithm can be jeopardized from a legal standpoint if it has disparate impacts on refugees from different origin countries (see Recital 71 of the EU GDPR [Vol22]). Second, fairness of employment outcomes across key subgroups has been identified by decision-makers as an important goal when designing algorithms. This sentiment is captured by Sjef van Grinsven, Director of Projects at COA:

> "*Because of the ever growing impact of AI on society, the topic of fairness is becoming increasingly important. For an organization such as COA, a semi-governmental organization working with a vulnerable target group, implementing fairness is not what you would call desirable: it's an absolute necessity. The challenge is to translate an abstract concept like fairness into concrete programmable choices.*"

Achieving the goal of incorporating group fairness into an algorithm for the dynamic refugee assignment problem is often not a straightforward task. First, there is a large and ever-growing number of ways that a policy maker may describe the fairness (or lack thereof) of an assignment of refugees to locations (see Section 1.2). Second, it is not clear how to extend fairness rules that are specified for static assignments to settings where refugees arrive and are assigned sequentially. To be practically useful to policy makers, any framework for incorporating group fairness into dynamic refugee resettlement must be *flexible* (seamlessly adapting to different fairness rules) and *simple* (not requiring the policy maker to reason about the intricacies of the dynamic arrivals of refugees).

1.1 Our contributions

In this paper, we introduce such a framework for incorporating group fairness into the dynamic refugee assignment problem (Section 2.2). Our framework does not require policy makers to define a notion of fairness that explicitly accounts for the stochastic arrival process of refugees. Instead, our framework only requires policy makers to specify an ex-post feasible fairness rule that, in turn, generates a *minimum requirement* on the average employment probability for each group. By only requiring the fairness rule to be defined when all refugee arrivals are known upfront, our framework is easy to understand by practitioners and thus desirable from an adoptability standpoint. At the same time, the minimum requirement can be generated using a variety of simple and interpretable definitions of group fairness from the literature, e.g., minmax, randomized, proportionally optimized within group. All together, our framework provides policy makers with a concrete way to evaluate whether a particular deployed algorithm achieves the specified fairness criteria (Section 2.3).

Equipped with our framework for reasoning about group fairness in the dynamic refugee resettlement setting, we perform a numerical analysis of how well existing approaches fare with respect to different fairness criteria (Section 5.2). We analyze data from two partner organizations in the US and the Netherlands and show that a natural Bid Price Control (BP), similar to [AGP$^+$21], can create frequent *unfair* group outcomes, when considering various fairness criteria considered in the literature. We also show that this is not simply a result of BP, but also occurs with a clairvoyant control that has advance knowledge of arrivals and assigns them optimally. Additionally, we show that a random assignment algorithm (a proxy for status quo procedures in the absence of optimization [BFH$^+$18]) yields outcomes that are inefficient in terms of total employment and routinely fails to achieve group fairness according to the rules considered.

Motivated by the finding that existing approaches can lead to unfair assignments on real data, we develop two algorithms with provable guarantees that hold for *any* ex-post feasible fairness rule.

- First, we propose a modification of BP (Section 3) that includes dual variables for each group. These dual variables amplify employment probabilities at the group level to help direct the most valuable locations towards the groups that need them the most in order to meet their minimum requirements. We refer to this algorithm as AMPLIFIED BID PRICE CONTROL (ABP). ABP has strong performance guarantees at the population level (Theorem 1) when compared to an offline solution that meets the minimum requirement of every group. Moreover, in our data-driven numerical study, it incurs only a minimal loss ($\approx 0-1\%$) in total employment when compared to BP and the offline problem without fairness constraints (see Table 2 in Section 5.3). This suggests that, in practice, incorporating group fairness considerations need not be costly in terms of total employment. On the theoretical side, ABP also provides strong performance guarantees at the group level for groups with reasonably large sizes (Theorem 2), such as age or educational level (see Table 1 in Section 5). However, for groups with small sizes, such as many based on country of origin in the US, the algorithm has no meaningful performance guarantees and can exhibit unfair outcomes in numerical studies (Section 5.3).
- Our second algorithm, dubbed CONSERVATIVE BID PRICE CONTROL (CBP), combines elements from ABP with *reserve capacity* at each location and occasional *greedy* decisions to overcome the poor performance of ABP with respect to small groups. Greedy steps sacrifice global efficiency to help boost the employment of groups that need it, and reserving capacity at each location helps ensure that regular assignments do not deplete the capacity of a location before a greedy decision can make use of it. With these changes, CBP has both population level performance guarantees (Theorem 3), and adds a guarantee that *all groups*, regardless of their size, will approximately meet their minimum requirements (Theorem 4). In our numerical results, we find that CBP incurs a small loss ($\approx 0-1\%$) in total employment when compared to ABP while approximately meeting the minimum requirements of all groups in all instances.

1.2 Related work

Refugee resettlement. In the context of refugee resettlement, recent work introduces data analytics and market design to improve operational efficiency. Our work is closest to the aforementioned works [BFH$^+$18, AAM$^+$21, BP22, AGP$^+$21] that optimize the assignment of refugees to locations, either in static or in dynamic settings, where the objective function is the average employment rate. [BP22, AGP$^+$21] use resolving techniques, backtest on historical data, and find that their online algorithms incur little loss relative to the optimal objective in hindsight. [BP22] incorporates an additional queuing/load balancing component in their objective to ensure that the resources at each location are evenly utilized throughout the year. The methods proposed by [BP22, AGP$^+$21] are currently being used in Switzerland and the US, respectively. [GP19] studies a variation wherein the employment score at a location is a submodular function of the assigned refugees. In terms of fairness, [AEM18] considers envy-freeness between locations. To the best of our knowledge, our work is the first to design algorithms that promote fairness for refugees. Closely connected to our motivation is [BM21], which identifies similar concerns to ours (Section 5.2) when studying a static refugee resettlement problem, but does so with a more restricted lens of group fairness and does not provide algorithmic solutions to overcome these concerns. Other works in refugee resettlement design mechanisms to satisfy preferences of refugees and locations [AE16, MR14, NNT21, JT18, ACGS18], or allow decision-makers to trade-off between refugees' preferences and employment maximization [ABH22, DKT$^+$19, OS22].

Methodological connections to online packing. Our formulation and algorithms are similar to those studied in online packing problems, including canonical quantity-based network revenue management [MVR99] and the AdWord problem [MSVV07]. These are sequential decision-making problems, in which a stream of T requests occurs, and a decision is made for each request: for network revenue management the decision is to *accept/reject*, for AdWords it is *where to accept* a request. In both problems accepting a request yields a reward (in AdWords, the reward depends on *where*), but also consumes a set of resources. The objective is to maximize the accumulated rewards and is often measured by regret, defined by the difference in accumulated rewards compared to the offline optimal solution, while satisfying capacity constraints [TVRVR04]. Classical bid price control [Wil92] and randomized LP approaches [TVR99] have long been known to achieve the fluid optimal $O(\sqrt{T})$ regret [TVR98] for these problems, though these approaches have been improved through resolving techniques, that eventually led to $O(1)$ regret guarantees [RW08, JK12, BW20, VB21, VBG21, BF19]. All of these approaches require resolving some LP from time to time. Compared with the aforementioned work, our framework includes fairness constraints which, after certain arrivals, may be impossible to meet at all. This is in contrast to usual packing problems where the decision maker can always satisfy capacity constraints by rejecting future requests. As a result, resolving techniques and the corresponding regret analyses are not immediately applicable to our setting.

An orthogonal set of approaches to online packing problems arises in online convex optimization [Haz16, AD15]. These approaches do not rely on resolving, but cannot easily adapt to the presence of sub-linear size groups, as we find in our work. Applying previous results, from either stream, leads to vacuous guarantees for small groups (see discussion after Theorem 2), forcing us to develop new algorithmic and analytical ideas to warrant good performance for all groups (Theorem 4).

Fairness considerations. Beyond refugee resettlement, fairness has been studied in many other operational settings such as kidney exchange [BFT13], food banks [SJBY22], online advertising [BCCM22], and pricing [CEL22]; see also [DAFST22] for a wider review of applications of algorithmic fairness in operations. We now expand on the connection to works that are closer to ours. The trade-off between optimizing a global objective and incorporating fairness considerations is often captured in the literature via the *price of fairness* [BFT11, BFT12]. Quantifying fairness considerations via a *requirement* for each group resembles a setting from the literature where one aims to minimize wait time in a queuing system subject to a requirement on the idle time of each server [AW10]. A major difference to these works is that our setting involves dynamic decisions occurring over a finite number of rounds; the uncertainty in arrivals introduces complexities, especially with respect to groups with small arrival probability. In contrast, the former line of works assumes a static optimization while the latter focuses on a steady-state behavior of the system.

The fairness rules that we consider in our work are closely related to two dynamic decision-making lines of work from the literature. The Proportionally Optimized fairness rule (Example 2) aims to mitigate the negative effect that the existence of a group may cause to members of another group by positing that the group should have performance at least as good as what it would have *in isolation*. This was initially studied in a stochastic online learning setting with disjoint groups

[RSWW18] and has been subsequently extended in a non-stationary setting with overlapping groups [BL20].[2] A key difference of our application is that individual cases are organically coupled via resource constraints; as a result, it is not clear how to define *in isolation*. To incorporate this aspect, we extend the rule by asserting that the average employment probability of each group is no less than what it would have been if the group optimized its assignments over hypothetical capacities proportional to its size (without considering other groups). The Maxmin fairness rule (Example 3) aims to optimize the performance of the group with the lowest average employment outcome; this is a classical fairness criterion [KS75, Raw04, KRT99, AS07, BFT11]. In a dynamic setting, there has been a sequence of recent works studying how to optimize this quantity [LIS14, MNR22, MXX22, MXX21]. A key difference in our setting is that we aim to optimize the global objective subject to a requirement on the lowest group's outcome, rather than maximizing the latter quantity directly.

Finally, fairness has been widely studied in the context of supervised machine learning with multiple different definitions mostly aiming to equalize a fairness metric across different groups starting from [HPS16, Cho17, KMR17, CG18]. The challenge in that line of literature is statistical: how can one try to use an inaccurate machine learning model? In contrast, we assume that the machine learning model is completely accurate (with respect to employment probabilities) and the challenge arises from the uncertainty in the arrival process.

2 Model

In this section, we define the dynamic refugee assignment problem (Section 2.1), introduce a framework that incorporates group fairness (Section 2.2), and formalize our objectives (Section 2.3).

2.1 The Dynamic Refugee Assignment Problem

A refugee resettlement agency or decision maker (DM) is faced with the task of assigning T refugee cases with labels $\{1, \dots, T\}$ to a set $\mathcal{M}$ of resettlement locations in a host country; we denote by M the number of locations. For simplicity, we assume that each refugee case consists of one individual.[3] Each location $j \in \mathcal{M}$ has a capacity s_j, i.e., it can have at most s_j cases assigned to it. We denote by $\hat{s}_j = s_j/T$ the fraction of cases that can be assigned to each location j. We refer to it as the *fractional capacity* of location j and denote the minimum fractional capacity by $\hat{s}_{\min} = \min_{j \in \mathcal{M}} \hat{s}_j$. The sequence of refugee cases is denoted by the stochastic process $\boldsymbol{\omega} \equiv (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_T)$, where $\boldsymbol{\theta}_t \in \Theta$ denotes the random feature vector associated with refugee case t. The feature vectors $\boldsymbol{\theta}_t$ are drawn identically and independently from a known probability distribution $\mathcal{P}$ over Θ and they include information that is known about the refugee such as their country of origin, gender, education level, etc. We denote the set of all possible realizations of this stochastic process by $\Omega \equiv \Theta^T$.

The feature vector $\boldsymbol{\theta}_t$ lets us extract the following relevant information about refugee case t. First, we let $w_{t,j} \equiv w_j(\boldsymbol{\theta}_t) \in [0, 1]$ denote the estimated probability that a refugee with features $\boldsymbol{\theta}_t$ finds employment at each location $j \in \mathcal{M}$ within a specified time frame of interest.[4] The function

[2] Since the groups can be overlapping, the latter work jointly optimizes the global objective as well as group objectives by positing that the whole population constitutes a group.

[3] Our results extend to cases that consist of multiple refugees if all members of each case belong to the same group. As we discuss in Section 6, handling intersectional groups is an interesting open direction.

[4] In the US, employment after 90 days of arrivals is the only integration outcome that is systematically tracked and reported. Many other host countries, for example the Netherlands and Switzerland, also track employment after arrival as one of the key integration metrics, although the time-frame of interest varies by country.

$\boldsymbol{w} : \Theta \to [0,1]^M$ is estimated by a supervised machine learning model and is available to the DM. We assume that $\mathcal{P}$ is such that $\boldsymbol{w}$ is drawn from a continuous cumulative distribution function over $[0,1]^M$, and often refer to $w_{t,j}$ as the score of case t at location j. Second, we let $g(t) \equiv g(\boldsymbol{\theta}_t) \in \mathcal{G}$ denote the (unique) group of refugee case t, where the set of all groups $\mathcal{G}$ has cardinality G. Group definitions are outside the purview of the DM and are decided by an external policy maker; for example, groups may be defined by level of education or country of origin.[5] We denote the expected proportion of cases that arrive from each group $g \in \mathcal{G}$ by $p_g \triangleq \mathbb{P}_{\boldsymbol{\theta}}(g(\boldsymbol{\theta}) = g)$.

Refugee cases arrive in ascending order of their labels and, at each period $t = 1, \ldots, T$, the DM observes the feature vector $\boldsymbol{\theta}_t$ associated with refugee case t and needs to decide a location for the case. This decision is non-anticipative and irrevocable: the DM needs to commit on the location for refugee case t prior to seeing the information of cases $t' > t$. Formally, we denote the DM's decision for refugee case t by the assignment vector $\mathbf{z}_t \equiv \mathbf{z}_t(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_t) \in \{0,1\}^M$, where $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_t$ is the information that has been revealed to the DM after the arrival of case t, and where the equality $z_{t,j} = 1$ holds if and only if refugee case t is assigned to location $j \in \mathcal{M}$. We say that a sequence of assignments is *feasible* if and only if each case is assigned to a single location in a way that does not violate any location's capacity; that is, a sequence of assignments $\mathbf{z}_1, \ldots, \mathbf{z}_T \in \{0,1\}^M$ is feasible if $(\mathbf{z}_1, \ldots, \mathbf{z}_T)$ is an element of the *fractional assignment polytope* formally defined as

$$\tilde{\mathcal{Z}} = \left\{ (\tilde{\mathbf{z}}_1, \ldots, \tilde{\mathbf{z}}_T) \in \mathbb{R}^{T \times M}_{\geq 0} : \quad \sum_{j \in \mathcal{M}} \tilde{\mathbf{z}}_{t,j} = 1 \quad \forall t \in \{1, \ldots, T\} \quad \text{and} \quad \sum_{t=1}^{T} \tilde{\mathbf{z}}_{t,j} \leq s_j \quad \forall j \in \mathcal{M}. \right\}$$

We denote the restriction of $\tilde{\mathcal{Z}}$ to the integers by $\mathcal{Z} = \tilde{\mathcal{Z}} \cap \{0,1\}^{T \times M}$.

In the absence of group fairness considerations, the DM aims to select a feasible sequence of integer assignments that achieves high average employment probability across refugee cases. Formally, a sequence of assignments $\mathbf{z}_1, \ldots, \mathbf{z}_T \in \mathcal{Z}$ induces a *global objective value* of $\frac{1}{T} \sum_{t=1}^{T} \sum_{j \in \mathcal{M}} w_{t,j} z_{t,j}$, where we recall that $w_{t,j}$ and $z_{t,j}$ are shorthand for $w_j(\boldsymbol{\theta}_t)$ and $z_{t,j}(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_t)$. To incorporate *group objectives*, we denote the set of cases that belong to group g, among the first t arrivals of realization $\boldsymbol{\omega}$, by $\mathcal{A}(g,t) \equiv \mathcal{A}(g,t,\boldsymbol{\omega}) \triangleq \{\tau \leq t : g(\boldsymbol{\theta}_\tau) = g\}$ and the cardinality of this set by $N(g,t) \equiv N(g,t,\boldsymbol{\omega}) \triangleq |\mathcal{A}(g,t,\boldsymbol{\omega})|$. A group is called *non-empty* for realization $\boldsymbol{\omega}$ if $N(g,T,\boldsymbol{\omega}) > 0$. This lets us define, for each group $g \in \mathcal{G}$, the average employment probability of that group as[6]

$$\alpha_g(\boldsymbol{\omega}) \triangleq \frac{1}{N(g,T,\boldsymbol{\omega}) \vee 1} \sum_{t \in \mathcal{A}(g,T)} \sum_{j \in \mathcal{M}} w_{t,j} z_{t,j}.$$

2.2 Our Framework for Group Fairness

Our goal is *not* to impose a definition of group fairness; rather, it is to provide flexible and expansive frameworks that allow policy makers to specify the group fairness desiderata they wish to attain. With this in mind, our framework allows external policy makers to specify a *minimum requirement* for every group and every realization of the dynamic refugee assignment problem. Formally, the policy maker selects a fairness rule $\mathcal{F}$ which provides a mapping $O_{g,\mathcal{F}} : \Omega \to [0,1]$ for each group $g \in \mathcal{G}$, where the quantity $O_{g,\mathcal{F}}(\boldsymbol{\omega})$ represents the minimum requirement on the average employment score of members of group g. Of course, not all requirements are achievable even when one is granted

[5] We emphasize the distinction between DM and policy maker to clarify what is in the purview of the algorithm and its designers. In particular, the definitions of groups and group fairness are exogenous to the algorithm.

[6] $\vee$ defines the maximum of two numbers, i.e., $a \vee b = \max\{a, b\}$, which defines the average as 0 for empty groups.

the benefit of hindsight. Hence, our framework necessitates that the provided requirements satisfy a notion of ex-post feasibility for a fractional version of the assignment problem.

Definition 1. *A fairness rule $\mathcal{F}$ is* ex-post feasible *if, for every sample path $\boldsymbol{\omega} \in \Omega$, there exist fractional ex-post decisions $\tilde{z}_{t,j}(\boldsymbol{\omega}) \in [0,1]$ for each case $t \in [T]$ and location $j \in \mathcal{M}$ such that*

$$(\tilde{\boldsymbol{z}}_1, \ldots, \tilde{\boldsymbol{z}}_T) \in \tilde{Z} \quad \textit{and} \quad \frac{1}{N(g,T) \vee 1} \sum_{t \in \mathcal{A}(g,T)} \sum_{j \in \mathcal{M}} w_{t,j} \tilde{z}_{t,j} \geq O_{g,\mathcal{F}}(\boldsymbol{\omega}) \quad \forall g \in \mathcal{G}.$$

Note that the constraints imply that $O_{g,\mathcal{F}}(\boldsymbol{\omega}) = 0$ for an empty group g, which we assume throughout.

As long as the specified fairness rule $\mathcal{F}$ satisfies ex-post feasibility, our algorithms will achieve favorable group performance guarantees despite not having the benefit of hindsight; see Sections 3 and 4. Of course, for this group fairness framework to be meaningful, it needs to capture existing fairness notions as special cases. To demonstrate the versatility of our framework, we provide three examples of ex-post feasible fairness rules. We note that for the first two examples to fulfill ex-post feasibility, it is necessary that Definition 1 allows for fractional assignments $\tilde{\mathbf{z}}$.

Example 1 (Random Fairness Rule)**.** *In many host countries, the status quo of refugee resettlement assignments is effectively random [BFH+ 18]. The random fairness rule protects groups from being worse off than this status quo benchmark, thus fulfilling the maxim of "primum non nocere" ("do no harm"). To mimic a random assignment, for each realization $\boldsymbol{\omega} \in \Omega$ and non-empty group $g \in \mathcal{G}$, we consider the fractional assignment $\tilde{z}^{\text{random}}_{t,j} = \frac{s_j}{\sum_{j' \in \mathcal{M}} s_{j'}}$ and set the minimum requirement as*

$$O_{g,\text{random}}(\boldsymbol{\omega}) = \frac{1}{N(g,T)} \sum_{t \in \mathcal{A}(g,T)} \sum_{j \in \mathcal{M}} w_{t,j} \tilde{z}^{\text{random}}_{t,j}.$$

Example 2 (Proportionally Optimized Fairness Rule)**.** *Though the random fairness rule fulfills "primum non nocere", it does not make use of synergies between different arrivals of one group. The proportionally optimized fairness rule posits that, for every realization $\boldsymbol{\omega} \in \Omega$, a non-empty group $g \in \mathcal{G}$ receives a capacity $s_{j,g} = N(g,T)\hat{s}_j$, proportional to the group size, for all locations $j \in \mathcal{M}$. It then optimizes their assignment over the fractional assignment polytope restricted to group g, i.e., $\tilde{\mathcal{Z}}_g = \{\tilde{\boldsymbol{z}} \in [0,1]^{\mathcal{A}(g,T) \times M} \colon \sum_{j \in \mathcal{M}} \tilde{z}_{t,j} \leq 1, \forall t \in \mathcal{A}(g,T);\ \sum_{t \in \mathcal{A}(g,T)} \tilde{z}_{t,j} \leq s_{j,g}, \forall j \in \mathcal{M}\}$. The resulting minimum requirement captures the maximum value a group would receive in isolation:*

$$O_{g,\text{pro}}(\boldsymbol{\omega}) = \max_{\tilde{\boldsymbol{z}} \in \tilde{\mathcal{Z}}_g} \frac{1}{N(g,T)} \sum_{t \in \mathcal{A}(g,T)} \sum_{j \in \mathcal{M}} w_{t,j} \tilde{z}_{t,j}.$$

This reflects the principle that a group should not be worse off due to the presence of other groups, which has been studied in sequential group fairness without resource constraints (see Section 1.2).

Example 3 (Max-Min Fairness Rule)**.** *As neither of the above examples guarantees a high minimum requirement for the most vulnerable (least employable) group, the well-studied max-min fairness rule is a natural alternative. The corresponding minimum requirement for realization $\omega \in \Omega$ is the same across all non-empty groups $g \in \mathcal{G}$ and maximizes the minimum average value across all groups:*

$$O_{g,\text{maxmin}}(\boldsymbol{\omega}) = \max_{\boldsymbol{z} \in \mathcal{Z}} \min_{g' \in \mathcal{G} \colon N(g',T) > 0} \frac{1}{N(g',T)} \sum_{t \in \mathcal{A}(g',T)} \sum_{j \in \mathcal{M}} w_{t,j} z_{t,j}.$$

Note that, although fairness rules are formally defined as a mapping from all possible realizations to minimum requirements, the examples above illustrate that fairness rules can be succinctly specified in an interpretable way.

Our work assumes that a policy maker has provided minimum requirements corresponding to an ex-post feasible fairness rule and we often drop the notational dependence on $\mathcal{F}$. In Section 5 we investigate the differences between the three rules we presented here. Finally, some of our results require the fairness rule $\mathcal{F}$ to satisfy a lower bound on a notion of slackness, defined as $\varepsilon_{g,\mathcal{F}} = \mathbb{E}_{\boldsymbol{\theta}}\left[\max_{j\in\mathcal{M}} w_j(\boldsymbol{\theta}) \,|\, g(\boldsymbol{\theta}) = g\right] - \mathbb{E}_{\boldsymbol{\omega}}[O_{g,\mathcal{F}}(\boldsymbol{\omega})]$, which is the expected difference between the maximum score conditional on a group g and its corresponding requirement. This quantity is non-negative by ex-post feasibility and captures the difficulty to achieve the minimum requirements.

Definition 2. *A fairness rule $\mathcal{F}$ has slackness ε if $\varepsilon_{g,\mathcal{F}} \geq \varepsilon > 0$ for every group g.*

Note that the only rule giving $\varepsilon_{g,\mathcal{F}} = 0$ must assign every case of group g to the location maximizing the employment score. In that sense, it is natural to assume that $\varepsilon_{g,\mathcal{F}}$ be greater than 0, though some of our results require it to be bounded away from 0.

2.3 Objectives for our Setting

The minimum requirements serve as an algorithmic benchmark: for a realization $\boldsymbol{\omega} \in \Omega$, we want an algorithm to achieve average employment probability for every group that is lower bounded by $O_{g,\mathcal{F}}(\boldsymbol{\omega})$. Although this is an ideal benchmark, which we use in our numerical results (Section 5), it is too strong to prove theoretical guarantees; in particular, there exists an instance for which no algorithm can be within 0.02 of the benchmark with high probability (see Appendix B.1). Hence, for our theoretical results, we consider a weaker benchmark that consists of the minimum of $O_{g,\mathcal{F}}(\boldsymbol{\omega})$ and its expectation over $\boldsymbol{\omega}' \sim \mathcal{P}^T$.[7] Formally, for a realization $\omega \in \Omega$, if an algorithm induces average employment probability $\alpha_g(\boldsymbol{\omega})$ for each group $g \in \mathcal{G}$, then its *g-regret* is

$$\mathcal{R}_{g,\mathcal{F}}(\boldsymbol{\omega}) \triangleq \min\left\{O_{g,\mathcal{F}}(\boldsymbol{\omega}), \mathbb{E}_{\boldsymbol{\omega}'}[O_{g,\mathcal{F}}(\boldsymbol{\omega}')]\right\} - \alpha_g(\boldsymbol{\omega}).$$

We now focus on our global objective, i.e., to maximize the average employment probability for all refugee cases independent of group. To evaluate how well an algorithm performs, we compare against the best sequence of fractional assignments for the realization $\boldsymbol{\omega}$. We note that, as before, we need to use fractional assignments to allow for general ex-post feasible fairness rules. Formally,

$$O^{\star}_{\mathcal{F}}(\boldsymbol{\omega}) \triangleq \max_{\tilde{\mathbf{z}}\in\tilde{\mathcal{Z}}} \frac{1}{T}\sum_{t=1}^{T}\sum_{j\in\mathcal{M}} w_{t,j}\tilde{z}_{t,j} \text{ s.t. } \forall g: \sum_{t\in\mathcal{A}(g,T)}\sum_{j\in\mathcal{M}} w_{t,j}\tilde{z}_{t,j} \geq N(g,T)\, O_{g,\mathcal{F}}(\boldsymbol{\omega}). \quad (\text{OFFLINE}_{\mathcal{F}})$$

Similarly, the global objective of an algorithm is to maximize the average employment probability across all refugee cases of the realization $\boldsymbol{\omega} \in \Omega$ subject to group fairness constraints. In particular, an algorithm aims to make sequential irrevocable integral decisions $z_t \in \{0,1\}^{\mathcal{M}}$ that minimize, with high probability, the *global regret*

$$\mathcal{R}_{\mathcal{F}}(\boldsymbol{\omega}) \triangleq \mathbb{E}_{\boldsymbol{\omega}'}[O^{\star}_{\mathcal{F}}(\boldsymbol{\omega}')] - \frac{1}{T}\sum_{t=1}^{T}\sum_{j\in\mathcal{M}} w_{t,j}z_{t,j},$$

[7] Note that $\mathbb{E}_{\boldsymbol{\omega}'}[O_{g,\mathcal{F}}(\boldsymbol{\omega}')]$ is also not a good benchmark by itself, as it may not be feasible for all realizations $\boldsymbol{\omega}'$.

while also achieving favorable g-regret guarantees for all groups $g \in \mathcal{G}$. In particular, we say that an algorithm has *favorable regret guarantees* if $\mathcal{R}_{\mathcal{F}}(\boldsymbol{\omega})$ and $\mathcal{R}_{g,\mathcal{F}}(\boldsymbol{\omega})$ $\forall g \in \mathcal{G}$ are upper bounded by a quantity that vanishes to zero as the number of cases T tends to infinity with high probability. Similar to above, we use $\mathbb{E}_{\boldsymbol{\omega}'}[O^{\star}_{\mathcal{F}}(\boldsymbol{\omega}')]$ instead of $O^{\star}_{\mathcal{F}}(\boldsymbol{\omega})$ as the benchmark because the latter is too strong to obtain a theoretical result. We construct an instance for which no algorithm can be within 0.02 of $O^{\star}_{\mathcal{F}}(\boldsymbol{\omega})$ and maintain low g-regret with high probability (see Appendix B.1).

Finally, for our regret to be well-defined, we require measurability of $O^{\star}_{\mathcal{F}}$ and $O_{g,\mathcal{F}}$ for every group $g \in \mathcal{G}$. We show in Appendix B.2 that the fairness rules we discussed before all satisfy this.

3 Amplified Bid Price Control

Our approach builds on the classical Bid Price Control that minimizes global regret (without fairness considerations) and was initially developed in the context of network revenue management [TVR98], which is an admission control problem with only accept/reject decisions. We show that a simple modification of this algorithm yields global regret that vanishes in T at a rate of $1/\sqrt{T}$ and g-regret that scales as $1/\sqrt{p_g \cdot T}$. In the next section, we extend our algorithm to attain g-regret that scales as $1/\sqrt{T}$.

Bid Price Control or BP as a shorthand aims to maximize the global objective subject to capacity constraints, i.e., find an assignment $\boldsymbol{z} \in \mathcal{Z}$ that maximizes $\sum_{t=1}^{T} \sum_{j \in \mathcal{M}} w_{t,j} z_{t,j}$. Myopically assigning each case to the location that maximizes its score may seem like a natural candidate algorithm to achieve this objective; yet, this quickly depletes the capacities at locations with universally good refugee employment prospects. As a result, optimally assigning a case needs to take into account not only the *present* case-location score but also the *opportunity cost* of a slot at location j, i.e., the potential decrease in future assignment scores due to depleted capacity at j. Bid Price Control quantifies the opportunity cost of each assignment and additively adjusts the employment probabilities by this amount. Effectively, this penalizes the selection of locations that are valuable in the future. That said, this adjustment of scores does not take into account group-fairness concerns and may severely violate minimum requirements (see Figure 4 of Section 5).

Our modification (Algorithm 1), which we term Amplified Bid Price Control or ABP as a shorthand, explicitly captures these fairness concerns by introducing group-dependent score amplifiers. In particular, the scores of each case are first scaled by their group's amplifier and are then additively adjusted by the opportunity cost of each potential assignment. These amplifiers help direct resources towards groups that need them the most in order to satisfy their minimum requirements. For example, a group that is likely to have a loose fairness constraint in $\text{OFFLINE}_{\mathcal{F}}$ would have a small amplifier and be assigned based on the tradeoff between $w_{t,j}$ and the locations' opportunity cost. To formalize the above intuition, we denote by $\mu^{\star}_j$ and $\lambda^{\star}_g$ the opportunity cost of location $j \in \mathcal{M}$ and the amplifier of group $g \in \mathcal{G}$, respectively. These quantities are initially instantiated (see line 1 in Algorithm 1); we elaborate on this instantiation below. ABP wants to assign case t to the location

$$J^{\text{WAB}}(t) \triangleq \arg\max_{j \in \mathcal{M}} (1 + \lambda^{\star}_{g(t)}) w_{t,j} - \mu^{\star}_j,$$

breaking ties in favor of the smallest index.[8] That said, if this location has no capacity, the final

[8]Note, however, that under our assumption of continuous scores, tie-breaking almost surely does not occur.

assignment $J^{\text{ABP}}(t)$ is a location with the most remaining capacity.

Algorithm 1: Amplified Bid Price Control

input : Capacity s_j for $j \in \mathcal{M}$; Time horizon T; Fairness rule $\mathcal{F}$ inducing scores $O_g(\omega)$

1 Set $(\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star) \leftarrow \arg\min_{\boldsymbol{\mu}\in\mathbb{R}_+^M, \boldsymbol{\lambda}\in\mathbb{R}_+^G} \mathbb{E}\left[L(\boldsymbol{\mu}, \boldsymbol{\lambda})\right]$ where $L(\boldsymbol{\mu}, \boldsymbol{\lambda})$ is defined in (LAGR)

2 **for** *case* $t = 1, \ldots, T$ *of group* $g(t) \in \mathcal{G}$ *with scores* $w_{t,j}$ *for* $j \in \mathcal{M}$ **do**

3 $\quad J^{\text{ABP}}(t) \leftarrow J^{\text{WAB}}(t)$ where $J^{\text{WAB}}(t) \leftarrow \arg\max_{j\in\mathcal{M}}(1 + \lambda^\star_{g(t)})w_{t,j} - \mu^\star_j$

4 $\quad$ **if** $J^{\text{WAB}}(t)$ *has no capacity* **then** $J^{\text{ABP}}(t) \leftarrow$ a location with most remaining capacity

We now expand on the instantiation of $\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star$. The problem without fairness constraints aims to find an assignment $\boldsymbol{z} \in \mathcal{Z}$ that maximizes $\sum_{t\in[T]}\sum_{j\in\mathcal{M}} w_{t,j}z_{t,j}$. The classical Bid Price Control computes opportunity costs $\boldsymbol{\mu}^\star$ by solving its Lagrangian relaxation over the fractional assignments without capacity constraints, i.e., the polytope $\widehat{\mathcal{Z}} \triangleq \{\widehat{\boldsymbol{z}} \in [0,1]^{T\times M} \colon \sum_{j\in\mathcal{M}} \widehat{z}_{t,j} = 1 \forall t\}$:

$$L(\boldsymbol{\mu}) = \frac{1}{T}\max_{\widehat{\boldsymbol{z}}\in\widehat{\mathcal{Z}}}\left(\sum_{t=1}^{T}\sum_{j\in\mathcal{M}} w_{t,j}\widehat{z}_{t,j} + \sum_{j\in\mathcal{M}}\mu_j\left(s_j - \sum_{t\in[T]}\widehat{z}_{t,j}\right)\right),$$

where $\mu_j \geq 0$ is the Lagrange multiplier associated with the capacity constraint for location j. For any vector $\boldsymbol{\mu} \geq 0$, $L(\boldsymbol{\mu})$ upper bounds the original objective.[9] BP selects Lagrange multipliers $\boldsymbol{\mu}^\star \in \arg\min_{\boldsymbol{\mu}\in\mathbb{R}_+^M}\mathbb{E}\left[L(\boldsymbol{\mu})\right]$ that give the tightest upper bound in expectation. For each case t, the optimal $\widehat{\boldsymbol{z}}$ in the inner maximization of $L(\boldsymbol{\mu}^\star)$ chooses the location j that maximizes $w_{t,j} - \mu_j^\star$. This matches the intuition that the multiplier $\mu_j^\star$ is the opportunity cost associated with location j.

We follow a similar approach while also incorporating group fairness constraints. In particular, denoting the Lagrange multiplier associated with the constraint for group g by λ_g, the Lagrangian relaxation $L(\boldsymbol{\mu}, \boldsymbol{\lambda})$ of our offline maximization program (OFFLINE$_\mathcal{F}$) is written as:

$$\max_{\widehat{\boldsymbol{z}}\in\widehat{\mathcal{Z}}}\frac{1}{T}\left(\sum_{t=1}^{T}\sum_{j\in\mathcal{M}} w_{t,j}\widehat{z}_{t,j} + \sum_{j\in\mathcal{M}}\mu_j\left(s_j - \sum_{t=1}^{T}\widehat{z}_{t,j}\right) + \sum_{g\in\mathcal{G}}\lambda_g\left(\sum_{t\in\mathcal{A}(g,T)}\sum_{j\in\mathcal{M}} w_{t,j}\widehat{z}_{t,j} - O_gN(g,T)\right)\right).$$

Rearranging terms, we find that $\widehat{z}_{t,j}$ appears as a coefficient for $(w_{t,j} - \mu_j + \lambda_{g(t)}w_{t,j}) = (1 + \lambda_{g(t)})w_{t,j} - \mu_j$. Similar to BP, the maximizing assignment $\widehat{\boldsymbol{z}} \in \widehat{\mathcal{Z}}$ now sets, for each case t, $\widehat{z}_{t,j} = 1$ for the location j that maximizes $(1 + \lambda_{g(t)})w_{t,j} - \mu_j$. Hence, the Lagrangian relaxation is

$$L(\boldsymbol{\mu}, \boldsymbol{\lambda}) = \frac{1}{T}\left(\sum_{t=1}^{T}\max_{j\in\mathcal{M}}((1 + \lambda_{g(t)})w_{t,j} - \mu_j) + \sum_{j\in\mathcal{M}}\mu_j s_j - \sum_{g\in\mathcal{G}}\lambda_g O_g N(g,T)\right). \tag{LAGR}$$

ABP picks opportunity costs $\boldsymbol{\mu}^\star$ and amplifiers $\boldsymbol{\lambda}^\star$ that minimize $\mathbb{E}_{\boldsymbol{\omega}}[L(\boldsymbol{\mu}, \boldsymbol{\lambda})]$ subject to non-negativity. As alluded to above, the scores of an arrival are amplified by λ_g before they are additively adjusted by opportunity cost μ_j. Further, by the same reasoning as for BP, $L(\boldsymbol{\mu}, \boldsymbol{\lambda}) \geq O^\star$ for any $\boldsymbol{\mu}, \boldsymbol{\lambda} \geq 0$. In particular, $L(\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star) \geq O^\star$ and thus $\mathbb{E}_{\boldsymbol{\omega}}[L(\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star)]$ is an upper bound for $\mathbb{E}_{\boldsymbol{\omega}}[O^\star]$. With $\mathbb{E}_{\boldsymbol{\omega}}[L(\boldsymbol{\mu}, \boldsymbol{\lambda})]$ being an expectation over convex functions (the maximum across linear functions), it is a convex function of $\boldsymbol{\mu}, \boldsymbol{\lambda}$ enabling the efficient instantiation of $\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star$ (line 1 of Algorithm 1).

In the following theorems, we establish global and g-regret guarantees for ABP.

[9] This is because every $\tilde{\boldsymbol{z}} \in \tilde{\mathcal{Z}}$ is also an element of $\widehat{\mathcal{Z}}$ as $\widehat{\mathcal{Z}}$ relaxes the capacity constraints; hence, when evaluating $\tilde{\boldsymbol{z}} \in \tilde{\mathcal{Z}}$ for the optimization problem in $L(\boldsymbol{\mu})$, this only adds a non-negative component to the original objective.

Theorem 1. *Fix an ex-post feasible fairness rule $\mathcal{F}$. For any $\delta > 0$,* Amplified Bid Price Control *has global regret $\mathcal{R}_{\mathcal{F}}^{\mathrm{ABP}} \leq \sqrt{\frac{\ln(M/\delta)}{T}} \cdot \left(\frac{1}{\hat{s}_{\min}} + 1\right)$ with probability at least $1 - \delta$.*

Theorem 2. *Fix an ex-post fairness rule $\mathcal{F}$. For any $\delta > 0$ and $g \in \mathcal{G}$,* Amplified Bid Price Control *has g-regret $\mathcal{R}_{g,\mathcal{F}}^{\mathrm{ABP}} \leq \sqrt{\frac{\ln(M/\delta)}{T}} \cdot \left(\frac{1}{\hat{s}_{\min}} + \sqrt{\frac{200}{p_g}}\right)$ with probability at least $1 - \delta$.*

We make two observations about the practicality of these bounds. First, the capacity of each location typically scales with the number of cases T;[10] which means that $\hat{s}_{\min}$ can be treated as a constant with respect to T. Therefore, the upper bound on global regret in Theorem 1 vanishes as T grows large with high probability, *e.g.*, when setting $\delta = 1/T$. Second, the upper bound on g-regret in Theorem 2 vanishes when a group's expected number of arrivals, $p_g \cdot T$, is large (*e.g.*, groups defined by age or education level; see Table 1). That said, this does not hold in all settings[11] and thus the bound may be vacuous. In Section 4 we design an algorithm to address this problem.

Our starting point towards proving the theorems is to analyze the global and group objectives when case t is always assigned to location $J^{\mathrm{WAB}}(t)$ irrespective of capacities. The objectives resulting from such an assignment have strong guarantees as we show in Lemma 3.1 and 3.2 using KKT conditions on $\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star$ and concentration arguments (see Appendix C.3 for full proofs).

Lemma 3.1. *For any t: $\mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)}\right] \geq \mathbb{E}\left[O^\star\right]$.*

Lemma 3.2. *For any t and g: $\mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)} \,\middle|\, g(t) = g\right] \geq \mathbb{E}\left[O_g\right] - \sqrt{\frac{1}{p_g T}}$.*

Of course, ignoring capacities is unlikely to yield a feasible solution since we have hard capacity constraints. However, until some location runs out of capacity, ABP follows exactly these assignments. As a result, we want to bound the first time when a location's capacity is depleted; we denote this time by T_{emp}. The following lemma shows that, with high probability, this event occurs in the last $\Delta^{\mathrm{WAB}} := \hat{s}_{\min}^{-1}\sqrt{\frac{1}{2}T\ln\left(\frac{M+2}{\delta}\right)}$ periods. The proof uses that, in expectation, J^{WAB} assigns at most s_j cases to each location j. Hence, by concentration arguments, ABP uses at most $\hat{s}_j t + O(\sqrt{t\ln(1/\delta)})$ capacity in a location j up to case t, and the algorithm does not deplete capacity in any location until assigning about $T - O(\sqrt{T\ln(1/\delta)})$ cases (see Appendix C.4 for a full proof).

Lemma 3.3. *For any $\delta > 0$ with probability at least $1 - \frac{M\delta}{M+2}$: $T_{\mathrm{emp}} \geq T - \Delta^{\mathrm{WAB}}$.*

The proofs of Theorems 1 and 2 follow from combining Lemma 3.3 with Lemmas 3.1 and 3.2 respectively and are provided in Appendix C.1 and C.2.

4 Conservative Bid Price Control

The poor performance of ABP for small groups (i.e., groups with small p_g) occurs due to two issues. First, ABP is designed to guarantee that the expected score of an arrival of a group is no less than the minimum requirement of that group. Roughly speaking, by central limit theorem, the average employment probability of a group concentrates near the expected score of one arrival of that group. That said, the rate at which this concentration occurs is inversely dependent on the size of

[10] *E.g.*, in Switzerland, this is government policy as assignments are made proportionally across regions [BFH+18].
[11] *E.g.*, in the U.S., out of 1,175 cases, 12 out of 26 groups defined by country of origin have a handful of cases.

the group. Second, consider an arrival t after T_{emp}, and suppose all the capacity at location $J^{\text{WAB}}(t)$ is depleted. The expected score of this arrival may be below the group's minimum requirement. For a group with few arrivals, this may have a significant effect on the group's average score.

Our second algorithm (Algorithm 2), termed CONSERVATIVE BID PRICE CONTROL, or CBP for short, is designed to address both of these issues by modifying ABP in two ways. First, CBP proactively offers high-score assignments to groups which are predicted to not meet their expected requirement otherwise. Second, CBP reserves capacity at locations so that late arrivals can still have favorable assignments. With these modifications, CBP approximately maintains the global guarantee of ABP (Theorem 3) while also obtaining strong guarantees for all groups (Theorem 4).

We first explain an intuition for how CBP predicts whether an arrival should receive a high-score assignment. Suppose for a group g a clairvoyant knew its future arrivals and could assign each arrival greedily to the score-maximizing location. If assigning t to $J^{\text{WAB}}(t)$ and all subsequent arrivals of group $g(t)$ greedily is insufficient for $g(t)$ to meet its requirement, then our clairvoyant would assign arrival t to the greedy location, denoted by $J^{\text{GR}}(t)$. In contrast, if it is sufficient, our clairvoyant would assign arrival t to $J^{\text{WAB}}(t)$. To formalize this intuition, let $J^{\text{CBP}}(\tau)$ denote the assignment that CBP made at time τ. Then, our clairvoyant condition is

$$\sum_{\tau \in \mathcal{A}(g,t-1)} w_{\tau, J^{\text{CBP}}(\tau)} + w_{t, J^{\text{WAB}}(t)} + \sum_{\tau \in \mathcal{A}(g,T) \setminus \mathcal{A}(g,t)} w_{\tau, J^{\text{GR}}(\tau)} \geq N(g,T)\mathbb{E}_{\boldsymbol{\omega}}[O_g(\boldsymbol{\omega})].$$

The terms on the left-hand side denote (i) the scores of already-assigned arrivals, (ii) the score of the current arrival t under J^{WAB}, and (iii) the cumulative scores of assigning the remaining arrivals of group g greedily. Since there are exactly $N(g,T)$ summands on the left, this can be rewritten as

$$\sum_{\tau \in \mathcal{A}(g,t-1)} \left(w_{\tau, J^{\text{CBP}}(\tau)} - \mathbb{E}_{\boldsymbol{\omega}}[O_g(\boldsymbol{\omega})]\right) + w_{t, J^{\text{WAB}}(t)} - \mathbb{E}_{\boldsymbol{\omega}}[O_g(\boldsymbol{\omega})] + \sum_{\tau \in \mathcal{A}(g,T) \setminus \mathcal{A}(g,t)} \left(w_{\tau, J^{\text{GR}}(\tau)} - \mathbb{E}_{\boldsymbol{\omega}}[O_g(\boldsymbol{\omega})]\right) \geq 0.$$

As we are not a clairvoyant, we cannot implement the above condition; we thus derive a proxy for the last sum, which is the only term that requires knowledge of future arrivals. Specifically, let$\beta > 0$ be a hyperparameter to capture how conservative the DM wants to be, which we later describe how to set. For a group g and a case t, we set $\Psi(g,t,\beta)$ so that, with high probability determined by β, $\Psi(g,t,\beta) \leq \sum_{\tau \in \mathcal{A}(g,T) \setminus \mathcal{A}(g,t)} \left(w_{\tau, J^{\text{GR}}(\tau)} - \mathbb{E}\left[O_g\right]\right)$. Though $\Psi(g,t,\beta)$ gives a valid lower bound on the sum, that sum itself may be an overestimate on the effect of future greedy assignments, since capacity at some locations will be entirely depleted towards the end of the horizon. To compensate for this, we add a buffer $\text{Buf}(\beta)$ to the right-hand side. Dropping the score $w_{t,J^{\text{WAB}}(t)}$ from the condition for simplicity and letting $V_g[t-1] = \sum_{\tau \in \mathcal{A}(g,t-1)} \left(w_{\tau, J^{\text{CBP}}(\tau)} - \mathbb{E}_{\boldsymbol{\omega}}[O_g(\boldsymbol{\omega})]\right)$ we can write

$$V_{g(t)}[t-1] - \mathbb{E}_{\boldsymbol{\omega}}[O_{g(t)}(\boldsymbol{\omega})] + \Psi(g(t),t,\beta) \geq \text{Buf}(\beta). \qquad \text{(predict-to-meet)}$$

As we alluded to above, CBP wants to assign case t to $J^{\text{WAB}}(t)$ if this condition holds and to $J^{\text{GR}}(t)$ otherwise; we denote this wanted assignment by $J^{\text{WCB}}(t)$ (see line 3 of Algorithm 2). However, like in ABP where exceptions are made for capacity depletion (line 4 in Algorithm 1), CBP must make several exceptions for different types of capacity. To address the second issue —unfavorable assignments for cases that arrive late in the horizon— CBP reserves a small portion of capacity, $\text{Res}(\beta)$ at each location j. This splits the capacity at location j into two components: free-to-use capacity f_j and reserved capacity r_j. CBP prioritizes free-to-use capacity; reserves can only be used for greedy selections. Intuitively, this allows CBP to execute greedy assignments for late cases without the risk that non-greedy assignments have already exhausted the corresponding capacity. However, this does not preclude the case where capacity has been depleted due to greedy assignments. To mitigate this concern, we limit the number of reserves that a group g can access to at most $\mathrm{Cap}(\beta)$; we denote by e_g the remaining reserves that may be accessed by group g. This helps ensure that reserves remain in the system for very late cases which allows CBP to execute greedy assignments for groups to meet their requirement. In our analysis we use $\boldsymbol{f}(t)$, $\boldsymbol{r}(t)$, and $\boldsymbol{e}(t)$ to denote the values of $\boldsymbol{f}$, $\boldsymbol{r}$, and $\boldsymbol{e}$ before the arrival of case t.

Given these different types of capacities, we now describe how CBP assigns case t to location $J^{\mathrm{CBP}}(t)$. If $J^{\mathrm{WCB}}(t)$ has free-to-use capacity, i.e., $f_{J^{\mathrm{WCB}}}(t) > 0$, then $J^{\mathrm{CBP}}(t) = J^{\mathrm{WCB}}(t)$ and CBP consumes one unit of that capacity. This is also the assignment when the following three conditions hold: (i) group $g(t)$ has not exhausted its allocation of reserves, i.e., $e_{g(t)} > 0$, (ii) there are remaining reserves at $J^{\mathrm{WCB}}(t)$, i.e., $r_{J^{\mathrm{WCB}}(t)}(t) > 0$, and (iii) CBP wants to make a greedy assignment, i.e., the (predict-to-meet) condition does not hold. Then, $J^{\mathrm{CBP}}(t) = J^{\mathrm{WCB}}(t)$ and CBP consumes one unit of the reserve. If neither of the previous conditions hold, CBP sets $J^{\mathrm{CBP}}(t)$ as a location with the most free-to-use capacity remaining. Finally, when all free-to-use capacities are depleted CBP gives up on reserves and relabels all capacities as free-to-use.

Algorithm 2: CONSERVATIVE BID PRICE CONTROL

input : Capacities s_j; Time horizon T; Fairness rule $\mathcal{F}$; Conservative hyperparameter β
1 Set $(\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star) \leftarrow \arg\min \mathbb{E}\left[L(\boldsymbol{\mu}, \boldsymbol{\lambda})\right]$ where $L(\boldsymbol{\mu}, \boldsymbol{\lambda})$ is defined in (LAGR)
2 Set $\{\Psi(g,t,\beta)\}_{g \in \mathcal{G}, t \in [T]}, \mathrm{Buf}(\beta), \mathrm{Cap}(\beta), \mathrm{Res}(\beta)$ by (PAR)
3 $e_g = \mathrm{Cap}(\beta), r_j = \mathrm{Res}(\beta);\ f_j = s_j - \mathrm{Res}(\beta), \forall g \in \mathcal{G}, j \in \mathcal{M}$
4 **for** $t = 1 \ldots T$ *of group* $g(t) \in \mathcal{G}$ *with scores* $w_{t,j}$ *for* $j \in \mathcal{M}$ **do**
5 | **if** (predict-to-meet) **then** $J^{\mathrm{WCB}}(t) \leftarrow J^{\mathrm{WAB}}(t)$ **else** $J^{\mathrm{WCB}}(t) \leftarrow J^{\mathrm{GR}}(t)$
6 | **if** $(f_{J^{\mathrm{WCB}}(t)} \geq 1)$ **then** $J^{\mathrm{CBP}}(t) \leftarrow J^{\mathrm{WCB}}(t);\ f_{J^{\mathrm{CBP}}(t)} \leftarrow f_{J^{\mathrm{CBP}}(t)} - 1$
7 | **elif** $(\min\{e_{g(t)}, r_{J^{\mathrm{WCB}}}\} > 0)$ **and** (**not** predict-to-meet) **then**
8 | $\quad J^{\mathrm{CBP}}(t) \leftarrow J^{\mathrm{WCB}}(t);\ r_{J^{\mathrm{CBP}}(t)} \leftarrow r_{J^{\mathrm{CBP}}(t)} - 1;\ e_{J^{\mathrm{CBP}}(t)} \leftarrow e_{J^{\mathrm{CBP}}(t)} - 1$
9 | **else** $J^{\mathrm{CBP}}(t) \leftarrow \arg\max_{j \in \mathcal{M}} f_j;\ f_j \leftarrow f_j - 1$
10 | **if** $\sum_{j \in \mathcal{M}} f_j = 0$ **then** $f_j \leftarrow r_j;\ r_j = 0, \forall j \in \mathcal{M}$

To finalize the presentation of the algorithm, we need to set its parameters (line 3 of Algorithm 2). In practice, these will be set in a data-driven way via Monte Carlo methods or parameter tuning (see Section 5 for further discussion). For our analysis, we instantiate them as:

$$\mathrm{Buf}(\beta) = 6\ln(1/\beta); \mathrm{Cap}(\beta) = \lceil 6\ln(1/\beta) \rceil; \mathrm{Res}(\beta) = G\mathrm{Cap}(\beta); \Psi(g,t,\beta) = \frac{-\ln(1/\beta)}{2\varepsilon_g}\ \forall g, t. \tag{PAR}$$

This requires that the fairness rule has positive slackness (Definition 2). A smaller value of β yields larger values for $\mathrm{Buf}(\beta)$, which makes it more difficult to satisfy (predict-to-meet) and results in higher reserves, leading the algorithm to be more conservative by assigning more cases greedily.[12]

In the following theorems, we establish global and g-regret guarantees for CBP.

Theorem 3. *Fix an ex-post feasible fairness rule $\mathcal{F}$ with slackness ε. For any $\delta > 0$, there exists T_0 s.t. $\forall T \geq T_0$,* CBP *with $\beta = \left(\frac{\delta}{12(M+G)T}\right)^{1/4}$ has $\mathcal{R}_{\mathcal{F}}^{\mathrm{CBP}} \leq \frac{16\ln(T/\delta)}{\hat{s}_{\min}\varepsilon}\sqrt{\frac{G}{T}}$ with probability at least $1 - \delta$.*

[12]This does not hold formally as a larger β need not always increase the number of greedy assignments.

Theorem 4. *Fix an ex-post feasible fairness rule $\mathcal{F}$ with slackness ε. For any $\delta > 0$ and $g \in \mathcal{G}$, there exists T_1 s.t. $\forall T \geq T_1$,* CBP *with $\beta = \left(\frac{\delta}{12(M+G)T}\right)^{1/4}$ has $\mathcal{R}^{\text{CBP}}_{g,\mathcal{F}} \leq \frac{200\ln(T/\delta)}{\hat{s}_{\min}\varepsilon}\sqrt{\frac{G}{T}}$ with probability at least $1-\delta$.*

CBP addresses the main shortcoming of ABP: Theorem 4 proves vanishing g-regret with no dependence on p_g. This comes at a cost: the bound in Theorem 3 is weaker than that in Theorem 1 in its dependence on G. This occurs due to the roughly $\sqrt{p_g T}/\varepsilon$ greedy steps that CBP may execute for each group g, each of which trades off a benefit in group outcomes for a loss in global objective. A weakness of Theorem 4 is that it only holds for large T.[13] That said, our numerical results show CBP performing well in practice, for both global objective and g-regret, including in settings where T is not large. We now provide a proof outline for Theorems 3 and 4.

Global regret. A vital quantity in the analysis of CBP is the first time a location runs out of free capacity. In a slight abuse of notation, we denote this quantity by T_{emp} as it plays the same role as in the analysis of ABP. The key difference in lower bounding T_{emp} for CBP arises from the greedy selections before T_{emp}. Because the number of greedy steps is small, we can bound T_{emp} similarly to Lemma 3.3 for ABP. Denoting $\Delta^{\text{WCB}} = \hat{s}_{\min}^{-1}\left(\frac{11\sqrt{GT}\ln(1/\beta)}{\varepsilon} + \frac{32G\ln(1/\beta)}{\varepsilon^2}\right)$, we show:

Lemma 4.1. *With probability at least $1-(2T+1)(M+G)\beta^4$, we have $T_{\text{emp}} \geq T - \Delta^{\text{WCB}}$.*

Proof sketch. We introduce a fictitious system in which each location has unlimited free-to-use capacity. In this fictitious system, the clause in line 6 is always triggered, so CBP assigns every case to $J^{\text{WCB}}(t)$. In the real system, for every $t \leq T_{\text{emp}}$, CBP, also assigns to $J^{\text{CBP}}(t) = J^{\text{WCB}}(t)$; hence CBP makes the same assignments in the fictitious and the real system when $t \leq T_{\text{emp}}$.

We then show that the total number of greedy steps in both the fictitious system and the real system are upper bounded by $\tilde{O}(\frac{\sqrt{GT}}{\varepsilon})$ with high probability (Lemma D.7). The intuition is as follows: By Lemma 3.2, the difference between expected requirement and the realized total score of a group g is around $\sqrt{p_g T}$. Therefore, to achieve condition predict-to-meet, there is a deficit of $\sqrt{p_g T}$ to be filled by greedy steps. By the slackness property (Definition 2), a greedy step (in expectation) fills at least ε of the gap toward the minimum requirement. Therefore, a group g takes around $\sqrt{p_g T}/\varepsilon$ greedy steps to cover the deficit. The Cauchy-Schwarz inequality then gives a bound on the total number of greedy steps of all groups. The formal proof relies on a concentration bound motivated by work on conservative bandits [WSLS16] and is provided in Appendix D.3. □

Proof of Theorem 3. The global average score can be bounded by the score obtained by T_{emp}, i.e., $\frac{1}{T}\sum_{t=1}^{T} w_{t,J^{\text{CBP}}(t)} \geq \frac{1}{T}\sum_{t=1}^{T_{\text{emp}}} w_{t,J^{\text{CBP}}(t)}$. For the first T_{emp} cases, with $J^{\text{CBP}}(t) = J^{\text{WCB}}(t)$, CBP assigns either to $J^{\text{WAB}}(t)$ or to the greedy assignment $J^{\text{GR}}(t)$. As a result,

$$w_{t,J^{\text{CBP}}(t)} \geq \min\left(w_{t,J^{\text{WAB}}(t)}, w_{t,J^{\text{GR}}(t)}\right) = w_{t,J^{\text{WAB}}}(t).$$

Then, the global objective of CBP is at least $\frac{1}{T}\sum_{t=1}^{T_{\text{emp}}} w_{t,J^{\text{WAB}}(t)} \geq \frac{1}{T}\sum_{t=1}^{T} w_{t,J^{\text{WAB}}(t)} - {}^{(T-T_{\text{emp}})}\!/_{T}$ where the last inequality comes from $w_{t,j} \in [0,1]$. With i.i.d. cases, by Hoeffding inequality (Fact 3):

$$\mathbb{P}\left\{\sum_{t=1}^{T} w_{t,J^{\text{WAB}}(t)} < T\mathbb{E}\left[w_{1,J^{\text{WAB}}(1)}\right] - \sqrt{2T\ln(1/\beta)}\right\} \leq \beta^4.$$

[13] This in part due to T_1 relying on δ, which is unavoidable for any algorithm as there are instances for which, for any $\delta > 0$, if $T < 0.15/\delta$, with probability δ the g-regret of a group is bounded below by a constant (Appendix D.2).

By Lemma 3.1, we have $\mathbb{E}\left[w_{1,J^{\mathrm{WAB}}(1)}\right] \geq \mathbb{E}[O^\star]$. As a result, with probability at least $1-\beta^4$, we have $\frac{1}{T}\sum_{t=1}^{T} w_{t,J^{\mathrm{WAB}}(t)} \geq \mathbb{E}[O^\star] - \sqrt{\frac{2\ln(1/\beta)}{T}}$. Lemma 4.1 shows that $T_{\mathrm{emp}} \geq T - \Delta^{\mathrm{WCB}}$ with probability at least $1-(2T+1)(M+G)\beta^4$. By union bound, the probability that the previous two events both occur is at least $1-(2T+1)(M+G)\beta^4 - \beta^4 \geq 1-(5T+3)(M+G+1)\beta^4$. Under these two events, we can lower bound the global average score

$$\frac{1}{T}\sum_{t=1}^{T} w_{t,J^{\mathrm{WAB}}(t)} - \frac{(T-T_{\mathrm{emp}})}{T} \geq \mathbb{E}[O^\star] - \sqrt{\frac{2\ln(1/\beta)}{T}} - \frac{\Delta^{\mathrm{WCB}}}{T}.$$

To get the desired result, focus on $\delta < 1$ (the result holds trivially for $\delta \geq 1$) and let

$$T_0 = \frac{12(M+G)}{\varepsilon^2}.$$

As $\beta = \left(\frac{\delta}{12(M+G)T}\right)^{1/4}$ and $T \geq T_0$, we have $\beta \leq e^{-0.5}$ which implies $\sqrt{2\ln(1/\beta)} < 2\ln(1/\beta)$. As a result, with probability at least $1-(5T+5)(M+G)\beta^4$, we have $\mathcal{R}_{\mathcal{F}}^{\mathrm{CBP}} \leq \frac{\Delta^{\mathrm{WCB}}}{T} + \frac{2\ln(1/\beta)}{\sqrt{T}}$. Moreover, for this β and $T \geq T_0$, we have $13\ln(1/\beta) \leq 4\ln(12(M+G)T/\delta) \leq 8\ln(T/\delta)$. Hence, with probability at least $1-\delta\frac{(5T+5)(M+G)}{12(M+G)T} \geq 1-\delta$, CBP has global regret

$$\mathcal{R}_{\mathcal{F}}^{\mathrm{CBP}} \leq \frac{13\ln(1/\beta)}{\hat{s}_{\min}\varepsilon}\sqrt{\frac{G}{T}}\left(1+\frac{3\sqrt{G}}{\varepsilon\sqrt{T}}\right) \leq \frac{8\ln(T/\delta)}{\hat{s}_{\min}\varepsilon}\sqrt{\frac{G}{T}}\left(1+\frac{3\sqrt{G}}{\varepsilon\sqrt{T}}\right). \tag{1}$$

Noting that $\frac{3\sqrt{G}}{\varepsilon\sqrt{T}} \leq 1$, this completes the proof. $\square$

Group regret. The crux of proving Theorem 4 is to connect the performance of a group $\alpha_g(\boldsymbol{\omega})$ with the condition (predict-to-meet). We want to show that, with high probability, if a group meets condition (predict-to-meet) in any period, then its outcome score α_g will end up near $\mathbb{E}[O_g]$. To formalize this, we define an event $\mathcal{S}_{\mathrm{BP}}$ such that conditioned on it, the following are true: (i) $T_{\mathrm{emp}} \geq T - \Delta^{\mathrm{WCB}}$, and (ii) if a group g fulfills (predict-to-meet) in some period t, where $g = g(t)$, then $\alpha_g \geq \mathbb{E}[O_g] - \frac{12\Delta^{\mathrm{WCB}}}{T}$. In the following Lemma we show that $\mathcal{S}_{\mathrm{BP}}$ occurs with high probability. Conditioned on this event, proving Theorem 4 then reduces to showing the following: (a) for any group that meets condition (predict-to-meet) in some period, the result guarantees a good outcome; (b) a group with many arrivals is unlikely to never meet condition (predict-to-meet) in any period, and (c) a group with few arrivals, under CBP has its cases assigned almost exclusively to their greedy location.

Lemma 4.2. *Suppose that $T \geq 36\Delta^{\mathrm{WCB}}$. Then $\mathbb{P}[\mathcal{S}_{\mathrm{BP}}] \geq 1-(3T+3)(M+G)\beta^4$.*

Proof sketch. Lemma 4.1 bounds the probability of $T_{\mathrm{emp}} \geq T - \Delta^{\mathrm{WCB}}$. To prove the lemma we need to show, with high probability, that if a group $g(t)$ fulfills condition (predict-to-meet) in some period t, then it ends up with $\alpha_g \geq \mathbb{E}[O_g] - \frac{12\Delta^{\mathrm{WCB}}}{T}$ at the end of the horizon. Consider the last period t in which group $g(t)$ fulfills condition (predict-to-meet). In every future period τ in which there is an arrival of group $g(t)$, the condition does not hold true. Thus, in those periods, $J^{\mathrm{WCB}}(\tau) = J^{\mathrm{GR}}(\tau)$. Moreover, we defined $\Psi(g,t,\beta)$ (see paragraph above (predict-to-meet)) to be a valid lower bound with high probability (Lemma D.8), so $\Psi(g,t,\beta) \leq \sum_{\tau \in \mathcal{A}(g,T)\setminus\mathcal{A}(g,t)} w_{\tau,J^{\mathrm{GR}}(\tau)}$. In particular, we then have $\sum_{\tau \in \mathcal{A}(g,T)} w_{\tau, J^{\mathrm{WCB}}(\tau)} \geq N(g,T)\mathbb{E}\left[O_g\right] + \mathrm{Buf}(\beta)$. For every case τ before T_{emp}, we have $J^{\mathrm{CBP}}(\tau) = J^{\mathrm{WCB}}(\tau)$. Therefore,

$$\sum_{\tau \in \mathcal{A}(g,T)} w_{\tau, J^{\mathrm{CBP}}(\tau)} \geq N(g,T)\mathbb{E}\left[O_g\right] + \mathrm{Buf}(\beta) - \left(N(g,T) - N(g,T_{\mathrm{emp}})\right).$$

Dividing both sides by $N(g,T)$ gives $\alpha_g \geq \mathbb{E}\left[O_g\right] + \frac{\mathrm{Buf}(\beta) - (N(g,T) - N(g,T_{\mathrm{emp}}))}{N(g,T)}$. Now for sufficiently large groups, we can show that $\frac{N(g,T) - N(g,T_{\mathrm{emp}})}{N(g,T)}$ is of the order of $\frac{T - T_{\mathrm{emp}}}{T}$ with high probability. For small groups, the choice of $\mathrm{Buf}(\beta)$ ensures that $\mathrm{Buf}(\beta) \geq N(g,T) - N(g,T_{\mathrm{emp}})$. Therefore, α_g is roughly lower bounded by $\mathbb{E}\left[O_g\right] - \frac{T - T_{\mathrm{emp}}}{T}$. The full proof is provided in Appendix D.4. □

We prove the theorem by using Lemma 4.2 to establish it separately for groups of sufficiently large size, and ones of very small size. Lemma 4.3 combines the slackness property and concentration bounds to prove the bound of Theorem 4 for groups g with large enough p_g. In particular, it shows that the total greedy scores of cases in group g that arrive before T_{emp} is large compared to the minimum requirement of group g. Therefore, condition (predict-to-meet) is met for at least one case (with high probability), allowing us to apply Lemma 4.2. The formal proof is in Appendix D.5.

Lemma 4.3. *Suppose that* $T \geq 36\Delta^{\mathrm{WCB}}$*. For a group* g*, if* $p_g \geq \frac{54 \ln^2(1/\beta)}{\varepsilon_g^2 T}$*, then* $\alpha_g \geq \mathbb{E}\left[O_g\right] - \frac{12\Delta^{\mathrm{WCB}}}{T}$ *with probability at least* $1 - (5T+5)(M+G)\beta^4$*.*

Now, consider a group g with small p_g. If condition (predict-to-meet) is satisfied for one of its cases, Lemma 4.2 holds and its average score is guaranteed to nearly exceed its expected minimum requirement with high probability. If, on the other hand, the predict-to-meet condition is not met for its cases before $T - T_{\mathrm{emp}}$, then they all receive greedy assignments. We then show that with high probability, its cases after $T - T_{\mathrm{emp}}$, but before free-to-use-capacities are exhausted (see Line 10 in CBP), all receive greedy assignments from reserved capacity. The average score is thus the highest possible and is at least the minimum requirement by ex-post feasibility of the fairness rule. These discussions show that for a small group g, its average score is at least close to $\min(O_g, \mathbb{E}\left[O_g\right])$ with high probability. We summarize the result as follows with the formal proof given in Appendix D.6.

Lemma 4.4. *Suppose that* $T \geq 36\Delta^{\mathrm{WCB}}$*. For every group* g *with* $p_g < \frac{54 \ln^2(1/\beta)}{\varepsilon^2 T}$*, we have with probability at least* $1 - (5T+5)(M+G)\beta^4 - \frac{378 MG \ln^3(1/\beta)}{\varepsilon^2 T}$ *that* $\alpha_g \geq \min\left(\mathbb{E}\left[O_g\right], O_g\right) - \frac{12\Delta^{\mathrm{WCB}}}{T}$*.*

The proof of Theorem 4 follows from combining Lemmas 4.3 and 4.4 (see Appendix D.1).

5 Numerical experiments

In this section we construct instances based on real-world data from the US and the Netherlands. We compare the minimum requirements under different fairness rules, show the limitations of status-quo approaches in achieving them, and study the empirical performance of our algorithms.

Data and instances. We use (de-identified) data from 2016 that cover free cases among (i) adult refugees that were resettled by one of the largest US resettlement agencies ($T = 1,175$), and

(ii) asylum seekers that were resettled in the Netherlands[14] ($T = 1,543$). In both countries, the outcome of interest is whether or not the refugee/asylum seeker found employment within a given time period (90 days in the US, 2 years in the Netherlands).[15] For each case t we apply the ML model of Bansak et al. [BFH$^+$18] to infer the employment probability $w_{t,j}$ at each location j.[16] Moreover, we set each capacity s_j as the number of cases that were assigned to location j in 2016.

We define three scenarios: NL-Age and NL-Edu where groups in the Netherlands are defined by either *age* or *education level*, and US-CoO where groups in the US are defined by *country of origin*. In the NL-Age scenario, age is segmented into 10 brackets, whereas cases in the NL-Edu scenario fall into one of seven groups based on educational attainment and degrees earned. When referring to the groups, we index them in increasing order of their size for US-CoO, and in increasing order of their age and education level for NL-Age and NL-Edu, respectively. As some cases consist of multiple individuals, we use the age/education level of the primary applicant to categorize each case into a group. These group definitions were chosen for both illustrative purposes and based on their importance to our partner organizations, though our approaches apply beyond these specific ones. Table 1 provides summary statistics on the arrivals and locations.

	T	M	G	Smallest group size	Largest group size
US-CoO	1,175	27	26	1	331
NL-Age	1,543	35	10	38	230
NL-Edu	1,543	35	7	30	428

Table 1: Summary statistics

Our numerical results are displayed over 50 bootstrapped samples $\boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_{50}$ drawn i.i.d. from the empirical distribution of 2016 arrivals. However, to illustrate the performance of our methods on real-world, non-stationary data, we also backtest them on 2016 arrivals in the order in which they occurred, with our algorithms only using information that was available at the start of 2016 (Appendix A.4). Implementation details of all online algorithms can be found in Appendix A.1.

5.1 Minimum requirements

We start by discussing the minimum requirements prescribed by the Random, Proportionally Optimized and Maxmin fairness rules (Examples 1-3 in Section 2) for our three scenarios (see Figure 2). The minimum requirement under the Random fairness rule (orange) is always below that of the Proportionally Optimized fairness rule (green). This holds by construction: each group receives the same (fractional) capacity at each location under both rules but the latter rule optimizes assignments for each group. Interestingly, due to higher employment rates among younger populations, in the NL-Age scenario (Figure 2a) the minimum requirements under both the Random and Proportionally Optimized fairness rule decrease as age increases. Furthermore, the gap between these two

[14] In the Netherlands, we specifically focus on the population of status holders who are granted residence permits and are assigned to a municipality through the regular housing procedure. We exclude some subsets of status holders who fall outside the scope of the objectives or for whom data are unreliable. First, we exclude status holders who fell under the 2019 Children's Pardon. Second, we exclude resettlers / relocants / asylum seekers covered by the EU-Turkey deal due to ambiguity in the data recorded on their registration.

[15] The period lengths and group definitions we consider stem from private communication with respective agencies.

[16] When cases consist of multiple individuals, we aggregate them by calculating the probability of at least one member of each case finding employment.

requirements decreases with the group's age. This suggests that there is more "room to optimize"—or more synergies between people and places—among younger populations. Conversely, for small groups in the US-CoO scenario with little room to optimize, the requirements under Proportionally Optimized and Random are often equal. For groups with larger index in Figure 2c, the gap between the two rules is more pronounced. For the Maxmin fairness rule (grey), the requirements are equal across groups and have no clear relationship to the requirements of the other fairness rules. For example, for groups 1-5 of the NL-Age scenario, the Maxmin requirement is the least constraining whereas for groups 9 and 10 it is the most constraining.

The monotonic relationship between the minimum requirements of Random and Prorportional Optimized has a clear implication for the group fairness and total employment we can expect: for any benchmark that does not take into account the fairness rule when making decisions, the outcomes will appear fairer under Random than under Proportionally Optimized; and for algorithms like ABP and CBP that try to meet group fairness constraints, the global objective will be higher under Random than under Proportionally Optimized. Because Maxmin does not have a monotonic relationship with the other fairness rules, we cannot predict the outcomes under this fairness rule.

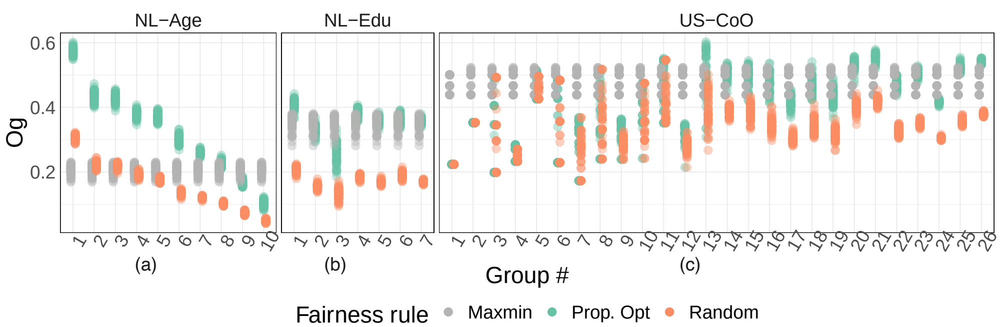

Figure 2: Minimum requirements for each group and scenario over bootstrapped arrival sequences.

5.2 Group fairness limitations of status quo approaches

We now discuss the group fairness of two algorithms that closely mimic status quo procedures: Rand and BP. Rand assigns each case randomly to a location with available capacity; this closely resembles status quo operations in the absence of optimization. BP (see Section 3) is similar to deployed algorithms (see [BP22, AGP+21]) and closely approximates the employment-maximizing (without fairness constraints) hindsight-optimal assignment, referred to as OPT. Although we discuss results in this section with respect to BP, similar insights arise for OPT (see Appendix A.2).

To measure group fairness, for a group g and a sample path $\boldsymbol{\omega}$, we define the *unfairness ratio* as

$$\mathrm{UFR}_g(\boldsymbol{\omega}) \triangleq \frac{O_{g,\mathcal{F}}(\boldsymbol{\omega}) - \alpha_g(\boldsymbol{\omega})}{O_{g,\mathcal{F}}(\boldsymbol{\omega})}.$$

This is an interpretable modification of g-regret: an unfairness ratio of 0.5 means that a group's average employment score is half of its minimum requirement prescribed by $\mathcal{F}$, and a negative unfairness ratio implies employment levels higher than the requirement. To aggregate across groups, we define the sample average maximum unfairness ratio as $\mathrm{UFR} = \frac{1}{50}\sum_{n=1}^{50} \max_{g\in\mathcal{G}} \mathrm{UFR}_g(\boldsymbol{\omega}_n)$.

We first examine the performance of Rand under our three fairness rules (see top of Figure 3). Recall that, for case t, their minimum requirement under the Random fairness rule is equivalent to their *expected* outcome under Rand. Rand, however, must choose a single location for each case. Thus, the unfairness ratios are generally centered around zero, but $\mathrm{UFR}_g(\boldsymbol{\omega})$ for a single sample path can be large, especially for the small groups in the US-CoO scenario. Results follow similar patterns under the Proportionally Optimized fairness rule but are more unfair. This arises because the minimum requirement obtained under the Proportionally Optimized fairness rule is larger than that obtained under the Random fairness rule (see Figure 2), especially for large groups.

We then consider the Maxmin fairness rule (see again top of Figure 2). In the NL-Age scenario, Rand is usually fair for groups 1-3, but not for 4-10 ($\mathrm{UFR}_g(\boldsymbol{\omega})$ of up to 0.8). The observation that unfairness increases with age is again a consequence of the higher overall employment rates among younger populations. To build intuition, consider a stylized example where younger groups are highly employable at any location, but older groups have only one location with high employment scores. In order to achieve their Maxmin requirement, a disproportionate number (relative to group size) of older cases would need to be assigned to that location. Since, under Rand, their expected allotment in each location is exactly proportional to their size, they do not meet their minimum requirement. In the NL-Edu and US-CoO scenarios, results under the Maxmin fairness rule reflect those obtained under proportionally optimized fairness.

Finally, we discuss the unfairness ratio of BP (bottom of Figure 3). For most groups (with the exception of certain small groups in the US-CoO scenario) BP is fair under the Random fairness rule. However, BP may be unfair for certain groups under both the Maxmin and Proportionally Optimized fairness rules. The unfairness ratio is particularly high under Maxmin fairness for groups 9 and 10 in the NL-Age scenario, group 3 in the NL-Edu scenario, and various groups in the US-CoO scenario. Therefore, although BP appears to be fairer than Rand for most groups, BP can be unfair, sometimes extremely so, for some groups.

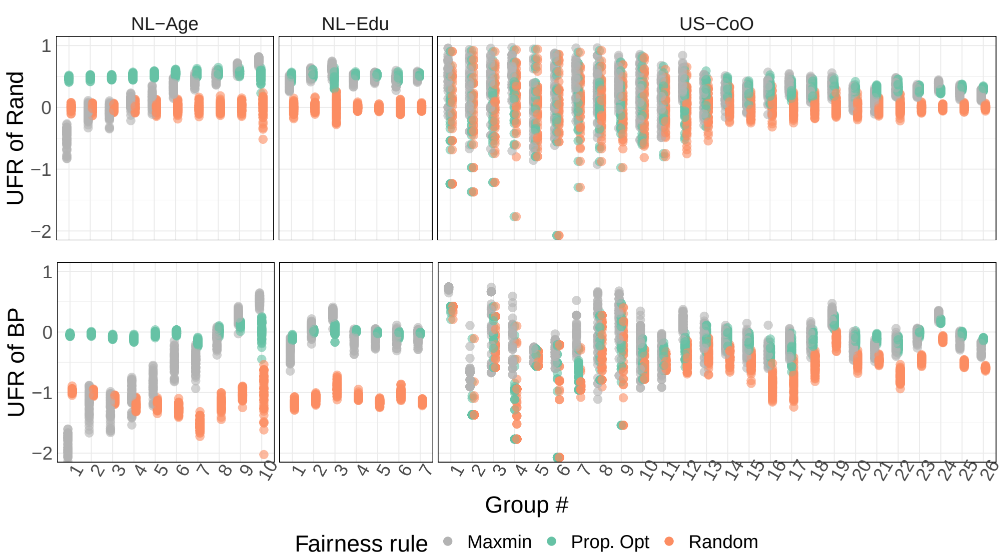

Figure 3: $\mathrm{UFR}_g(\boldsymbol{\omega})$ of Rand and BP with respect to the three fairness rules in each scenario.

5.3 Group and global objectives of our algorithms

We now compare the unfairness ratio and total employment score achieved by our algorithms (ABP and CBP) with the above status-quo approaches (Rand and BP) and the optimal offline benchmarks, OFFLINE$_{\mathcal{F}}$ and OPT. We note that the latter rely on knowledge of the sample path $\boldsymbol{\omega}$ and thus are not implementable in practice but are useful as upper bounds for the total employment with and without fairness constraints, respectively.

Group fairness. Our analysis finds that each of our benchmarks[17] exhibits a UFR of at least 50% in multiple instances (see Figure 4). In comparison, ABP and/or CBP improve upon UFR in all instances where it is positive, often dramatically, with the exception of the NL-Edu scenario under the Proportionally Optimized fairness rule. In this instance, OPT already has very small UFR, and therefore BP, ABP, and CBP perform similarly, and are all less fair than OPT due to the efficiency loss from making decisions online. Focusing on the scenarios from the Netherlands (which have only large groups, see Table 1), ABP has a maximum UFR of 9%; in the US-CoO scenario that does contain small groups, CBP has a maximum UFR of 8%. More generally, in line with our expectations from Theorems 2 and 4, ABP and CBP perform similarly well on UFR when all groups are large, and it is only in the US-CoO scenario that CBP does significantly better.

The existence of small groups in the US-CoO scenario warrants a more detailed investigation of each algorithm's UFR. Focusing on small (less than 10 arrivals) and large groups (more than 100), we find that CBP dramatically reduces the proportion of sample paths with positive UFR$_g(\boldsymbol{\omega})$ for small groups. Among large groups, CBP results in a larger proportion of sample paths with positive UFR$_g(\boldsymbol{\omega})$ than ABP, though the magnitude is small in these cases (see Table 3 in Appendix A.3 for more details). This demonstrates a trade-off: ABP guarantees group-fairness for large groups, whereas CBP introduces inefficiencies to help all groups meet their minimum requirement.

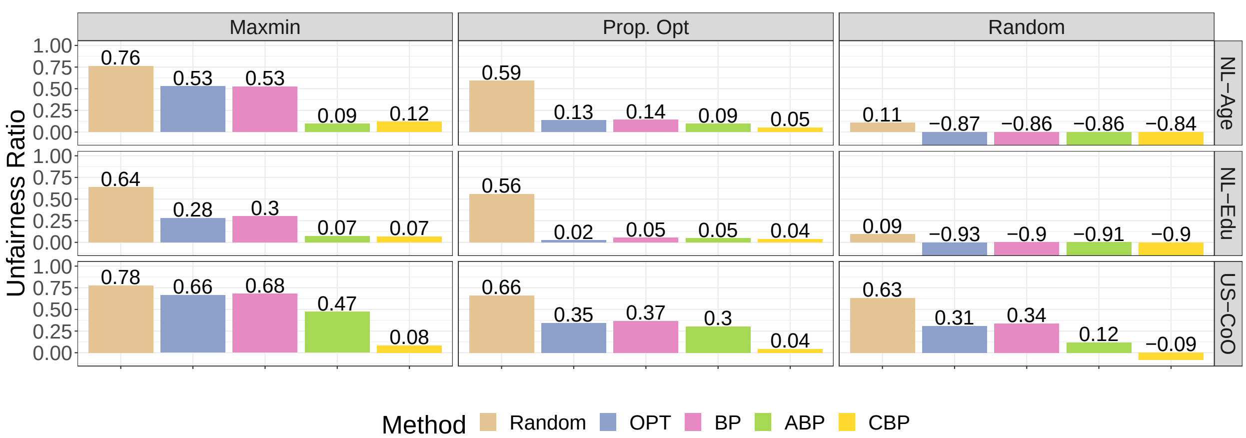

Figure 4: Average unfairness ratio under the three fairness rules for each scenario, aggregated across 50 bootstrapped sample paths. To increase clarity, negative values are truncated at -0.1.

Global objective. The last goal of our data-driven exploration is to compare the global objective of the approaches and benchmarks (see Table 2). By comparing OFFLINE$_{\mathcal{F}}$ to OPT we can characterize the loss in total employment necessary to accommodate group fairness. With just one exception, OFFLINE$_{\mathcal{F}}$ always obtains at least 99% of OPT, showing that group fairness may come at a low cost. In some cases, this occurs because the minimum requirements are met "for free", i.e., the fairness constraints are naturally fulfilled when maximizing total employment (see Figure 4). A second and more nuanced reason relates to patterns in the employment score vectors. When these are highly correlated, yet slightly offset, swapping the assignment of two cases in different groups

[17]This is with the exception of OFFLINE$_{\mathcal{F}}$ that has non-positive UFR by definition.

may yield small changes in total employment but large changes in average group employment.[18]

All of the online methods (BP, ABP, and CBP) average at least 95% of the total employment level under OPT. The ordering between these three is unsurprising since BP optimizes solely for total employment while ABP and CBP introduce inefficiencies with the goal of increasing degrees of fairness. However, ABP loses no more than 2% efficiency compared to BP, and CBP loses no more than 1% efficiency compared to ABP. In the NL-Edu scenario under the Proportionally Optimized fairness rule, where OPT results in low UFR, ABP and CBP not only result in the same UFR as BP (discussed above), but also produce the same total employment (Table 2). Therefore, as desired, when fairness is not a concern our algorithms revert back to maximizing total employment. In other scenarios, these small losses in efficiency should be considered alongside the improvement in fairness that they achieve. Indeed, for US-CoO under the random fairness rule, CBP and ABP are as efficient as BP while significantly reducing unfairness ratios.

	NL-Age			NL-Edu			US-CoO		
Rule:	Maxmin	PrOpt	Random	Maxmin	PrOpt	Random	Maxmin	PrOpt	Random
Rand	0.46	0.46	0.46	0.46	0.46	0.46	0.67	0.67	0.67
CBP	0.95	0.97	0.98	0.97	0.97	0.98	0.97	0.98	0.99
ABP	0.96	0.98	0.98	0.98	0.98	0.98	0.98	0.99	0.99
BP	0.98	0.98	0.98	0.98	0.98	0.98	0.99	0.99	0.99
OFFLINE$_{\mathcal{F}}$	0.97	1.00	1.00	0.99	1.00	1.00	0.99	1.00	1.00

Table 2: Average ratio of total employment achieved under each algorithm compared to OPT.

6 Conclusions

In this paper, we introduced the first formal approach to incorporating group fairness into the dynamic refugee assignment problem. Our approach explicitly addresses the practical needs of refugee resettlement stakeholders by offering flexibility to policy makers when specifying the notions of group fairness desiderata they wish to attain. Moreover, we developed new assignment algorithms that are guaranteed to achieve strong global and group-level performance guarantees for *any* feasible definition of group fairness chosen by the policy maker. Finally, through extensive numerical experiments on multiple real-world data, we showed that our online algorithms can achieve desired group-level fairness outcomes with very small changes in global performance.

Both our theoretical and empirical results demonstrate the performance of our two proposed algorithms, ABP and CBP. While BP achieves a higher overall average employment rate than our proposed methods, this results in unfair outcomes for certain groups under various fairness rules. In some settings, BP gets “lucky” and achieves group fairness, for example when the minimum requirements are quite loose as in the NL-Edu scenario under the Random fairness rule. ABP improves upon BP in terms of fairness, working well when all groups are relatively large and obtaining an overall average employment score within 2% of BP on real-world data. CBP sacrifices some efficiency ($\sim 1\%$) relative to ABP in order to achieve fairer outcomes for all groups regardless of size. Therefore, it is a natural choice in settings with small groups.

Our work opens up a number of promising directions for future research:

[18] For example, suppose there are just two locations, and members of Groups A and B have employment score vectors of $[.7, .4]$ and $[.68, .39]$ respectively. Then, swapping assignments of members of Group A and B changes total employment score by .01, but may improve UFR by much more.

- In this paper, we make the assumption that each individual belongs in a unique group; this makes sense when the group definition is based on a single attribute (country of origin, age, educational level, gender) and is important for our algorithms to know how to assign reserves. That said, in practice, refugees and asylum seekers have multiple attributes and their profile is defined as an intersection of the respective subgroups. Designing algorithms that account for such *intersectional* groups is a fundamental direction for future work.
- In this paper, we assume that a case consists of a single individual and our results can easily extend to the case where cases have multiple individuals who all belong to the same group. That said, many families consist of members of different gender, educational level, age, etc. Since families must be assigned to the same location, our methods do not readily apply to settings with within-case group-heterogeneity and this is an important extension.
- Although our framework is flexible enough to incorporate various fairness rules, these rules need to be defined as a minimum requirement on the average performance on the group. This precludes fairness rules that aim to optimize, say, for the performance achieved by the 10th-percentile of the group, which may be a meaningful statistic as it disallows a few good outcomes to outweigh the experience of the majority of the group members. Providing results that account for more general fairness rules can thus be a useful direction for future work.

Finally, we hope that future work may extend our algorithmic techniques for achieving group fairness that is specified by an *ex-post minimum requirement* (Section 2.2) to various other public policy settings in which individuals must be assigned dynamically (e.g., in health care and housing). More broadly, this work is a concrete example of how AI can be used to achieve complex and socially beneficial policy goals—an area with significant opportunity for future research.

Acknowledgement

We acknowledge funding from the Charles Koch Foundation, Google.org, Stanford Impact Labs, the National Science Foundation, and J-PAL for this work. The funders had no role in the data collection, analysis, decision to publish, or preparation of the manuscript.

References

[AAM+21] Narges Ahani, Tommy Andersson, Alessandro Martinello, Alexander Teytelboym, and Andrew C. Trapp. Placement optimization in refugee resettlement. *Oper. Res.*, 69(5):1468–1486, 2021. 2, 4

[ABH22] Avidit Acharya, Kirk Bansak, and Jens Hainmueller. Combining outcome-based and preference-based matching: A constrained priority mechanism. *Political Analysis*, 30(1):89–112, 2022. 5

[ACGS18] Haris Aziz, Jiayin Chen, Serge Gaspers, and Zhaohong Sun. Stability and pareto optimality in refugee allocation matchings. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 964–972, 2018. 5

[AD15] Shipra Agrawal and Nikhil R Devanur. Fast algorithms for online stochastic convex programming. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 1405–1424. SIAM, 2015. 5

[AE16] Tommy Andersson and Lars Ehlers. Assigning refugees to landlords in sweden: Stable maximum matchings. *Cahier de recherche*, (2016-08), 2016. 5

[AEM18] Tommy Andersson, Lars Ehlers, and Alessandro Martinello. Dynamic refugee matching. *Cahier de recherche*, (2018-10), 2018. 4

[AGP+21] Narges Ahani, Paul Gölz, Ariel D Procaccia, Alexander Teytelboym, and Andrew C Trapp. Dynamic placement in refugee resettlement. *arXiv preprint arXiv:2105.14388*, 2021. 2, 3, 4, 19

[AS07] Arash Asadpour and Amin Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 114–121, 2007. 6

[AW10] Mor Armony and Amy R Ward. Fair dynamic routing in large-scale heterogeneous-server systems. *Operations Research*, 58(3):624–637, 2010. 5

[BCCM22] Mohammad Hossein Bateni, Yiwei Chen, Dragos Florin Ciocan, and Vahab S. Mirrokni. Fair resource allocation in a volatile marketplace. *Oper. Res.*, 70(1):288–308, 2022. 5

[BF19] Siddhartha Banerjee and Daniel Freund. Good prophets know when the end is near. *Available at SSRN 3479189*, 2019. 5

[BFH+18] Kirk Bansak, Jeremy Ferwerda, Jens Hainmueller, Andrea Dillon, Dominik Hangartner, Duncan Lawrence, and Jeremy Weinstein. Improving refugee integration through data-driven algorithmic assignment. *Science*, 359(6373):325–329, 2018. 2, 3, 4, 8, 12, 18

[BFT11] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. The price of fairness. *Operations research*, 59(1):17–31, 2011. 5, 6

[BFT12] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. On the efficiency-fairness trade-off. *Management Science*, 58(12):2234–2250, 2012. 5

[BFT13] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013. 5

[BL20] Avrim Blum and Thodoris Lykouris. Advancing subgroup fairness via sleeping experts. In *11th Innovations in Theoretical Computer Science Conference*, 2020. 6

[BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013. 52

[BM21] Kirk Bansak and Linna Martén. Algorithmic decision-making, fairness, and the distribution of impact: Application to refugee matching in sweden, 2021. 4

[BP22] Kirk Bansak and Elisabeth Paulson. Outcome-driven dynamic refugee assignment with allocation balancing. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 1182–1183, 2022. 2, 4, 19

[BW20] Pornpawee Bumpensanti and He Wang. A re-solving heuristic with uniformly bounded loss for network revenue management. *Management Science*, 66(7):2993–3009, 2020. 5

[CEL22] Maxime C Cohen, Adam N Elmachtoub, and Xiao Lei. Price discrimination with fairness constraints. *Management Science*, 2022. 5

[CG18] Sam Corbett-Davies and Sharad Goel. The measure and mismeasure of fairness: A critical review of fair machine learning. *CoRR*, abs/1808.00023, 2018. 6

[Cho17] Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2):153–163, 2017. 6

[DAFST22] Maria De-Arteaga, Stefan Feuerriegel, and Maytal Saar-Tsechansky. Algorithmic fairness in business analytics: Directions for research and practice. *Production and Operations Management*, 31(10):3749–3770, 2022. 5

[DKT+19] David Delacrétaz, Scott Duke Kominers, Alexander Teytelboym, et al. Matching mechanisms for refugee resettlement. Technical report, University of Oxford, Oxford, UK, 2019. 5

[Dur19] Rick Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019. 34

[GP19] Paul Gölz and Ariel D. Procaccia. Migration as submodular optimization. In *The Thirty-Third Conference on Artificial Intelligence*, pages 549–556. AAAI Press, 2019. 4

[Haz16] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016. 5

[HPS16] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016. 6

[JK12] Stefanus Jasin and Sunil Kumar. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research*, 37(2):313–345, 2012. 5

[JT18] Will Jones and Alexander Teytelboym. The local refugee match: Aligning refugees' preferences with the capacities and priorities of localities. *Journal of Refugee Studies*, 31(2):152–178, 2018. 5

[KMR17] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, volume 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 43:1–43:23. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017. 6

[KRT99] Jon Kleinberg, Yuval Rabani, and Éva Tardos. Fairness in routing and load balancing. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 568–578. IEEE, 1999. 6

[KS75] Ehud Kalai and Meir Smorodinsky. Other solutions to nash's bargaining problem. *Econometrica: Journal of the Econometric Society*, pages 513–518, 1975. 6

[LIS14] Robert W Lien, Seyed MR Iravani, and Karen R Smilowitz. Sequential resource allocation for nonprofit operations. *Operations Research*, 62(2):301–317, 2014. 6

[Mar75] DH Martin. On the continuity of the maximum in parametric linear programming. *Journal of Optimization Theory and Applications*, 17(3):205–210, 1975. 36

[MNR22] Vahideh Manshadi, Rad Niazadeh, and Scott Rodilitz. Fair dynamic rationing. *Management science*, 2022. Forthcoming. 6

[MR14] Jesús Fernández-Huertas Moraga and Hillel Rapoport. Tradable immigration quotas. *Journal of Public Economics*, 115:94–108, 2014. 5

[MSVV07] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5):22–es, 2007. 5

[MVR99] Jeffrey I McGill and Garrett J Van Ryzin. Revenue management: Research overview and prospects. *Transportation science*, 33(2):233–256, 1999. 5

[MXX21] Will Ma, Pan Xu, and Yifan Xu. Fairness maximization among offline agents in online-matching markets. *arXiv preprint arXiv:2109.08934*, 2021. 6

[MXX22] Will Ma, Pan Xu, and Yifan Xu. Group-level fairness maximization in online bipartite matching. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 1687–1689, 2022. 6

[NNT21] Hai Nguyen, Thành Nguyen, and Alexander Teytelboym. Stability in matching markets with complex constraints. *Management Science*, 67(12):7438–7454, 2021. 5

[OS22] Nils Olberg and Sven Seuken. Enabling trade-offs in machine learning-based matching for refugee resettlement. *arXiv preprint arXiv:2203.16176*, 2022. 5

[Raw04] John Rawls. A theory of justice. In *Ethics*, pages 229–234. Routledge, 2004. 6

[RSWW18] Manish Raghavan, Aleksandrs Slivkins, Jennifer Vaughan Wortman, and Zhiwei Steven Wu. The externalities of exploration and how data diversity helps exploitation. In *Proceedings of the 31st Conference On Learning Theory*, pages 1724–1738, 2018. 6

[RW08] Martin I Reiman and Qiong Wang. An asymptotically optimal policy for a quantity-based network revenue management problem. *Mathematics of Operations Research*, 33(2):257–282, 2008. 5

[SJBY22] Sean R Sinclair, Gauri Jain, Siddhartha Banerjee, and Christina Lee Yu. Sequential fair allocation: Achieving the optimal envy-efficiency tradeoff curve. *Operations Research*, 2022. Forthcoming. 5

[TVR98] Kalyan Talluri and Garrett Van Ryzin. An analysis of bid-price controls for network revenue management. *Management science*, 44(11-part-1):1577–1593, 1998. 5, 10, 39

[TVR99] Kalyan Talluri and Garrett Van Ryzin. A randomized linear programming method for computing network bid prices. *Transportation science*, 33(2):207–216, 1999. 5

[TVRVR04] Kalyan T Talluri, Garrett Van Ryzin, and Garrett Van Ryzin. *The theory and practice of revenue management*, volume 1. Springer, 2004. 5

[UNH22] UNHCR. UNHCR Projected Global Resettlement Needs 2023. Technical report, 2022. 1

[VB21] Alberto Vera and Siddhartha Banerjee. The bayesian prophet: A low-regret framework for online decision making. *Management Science*, 67(3):1368–1391, 2021. 5

[VBG21] Alberto Vera, Siddhartha Banerjee, and Itai Gurvich. Online allocation and pricing: Constant regret via bellman inequalities. *Operations Research*, 69(3):821–840, 2021. 5

[Vol22] Nicholas Vollmer. Recital 71 eu general data protection regulation (eu-gdpr), Aug 2022. 3

[Wil92] Elizabeth Louise Williamson. *Airline network seat inventory control: Methodologies and revenue impacts*. PhD thesis, Massachusetts Institute of Technology, 1992. 5

[WSLS16] Yifan Wu, Roshan Shariff, Tor Lattimore, and Csaba Szepesvári. Conservative bandits. In Maria-Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of the 33nd International Conference on Machine Learning*, volume 48, pages 1254–1262, 2016. 15

A Supplementary material for Section 5

A.1 Algorithmic set-up

This section describes the details of the implementations of all online methods—BP, ABP, and CBP— used in Section 5.

ABP and CBP. As described in Section 3, the parameters $\boldsymbol{\lambda}^*$ and $\boldsymbol{\mu}^*$ used in ABP and CBP (Algorithms 1 and 2) are found by solving $\min_{\boldsymbol{\lambda}>0,\boldsymbol{\mu}>0} \mathbb{E}\left[L(\boldsymbol{\mu},\boldsymbol{\lambda})\right]$ where

$$L(\boldsymbol{\mu},\boldsymbol{\lambda}) = \sum_{t=1}^{T} \max_{j\in\mathcal{M}}((1+\lambda_{g(t)})w_{t,j} - \mu_j) + \sum_{j\in\mathcal{M}} \mu_j s_j - \sum_{g\in\mathcal{G}} \lambda_g O_g N(g,T).$$

Let the superscript k index a sample path of arrivals. The minimization problem is approximated by solving the following sample-average approximation across K sample paths:

$$\begin{aligned} \min_{\boldsymbol{\lambda}>0,\boldsymbol{\mu}>0} & \sum_{k=1}^{K} \left(\sum_{t=1}^{T} z_t^k + \sum_{j\in\mathcal{M}} \mu_j s_j - \sum_{g\in\mathcal{G}} \lambda_g O_g^k N^k(g,T) \right) \\ \text{s.t. } & z_t^k \geq ((1+\lambda_{g^k(t)})w_{t,j}^k - \mu_j) \ \forall j \in \mathcal{M},\ t \in \{1,...,T\} \end{aligned} \tag{2}$$

where O_g^k and $N^k(g,T)$ can be pre-computed offline for each sample path of arrivals and then treated as constants in Problem 2. Using the auxiliary variables z_t^k, Problem 2 is a linear program and thus computationally tractable.

After calculating $\boldsymbol{\lambda}^*$ and $\boldsymbol{\mu}^*$, Algorithm 1 can be implemented as described in Section 3. For Algorithm 2, we similarly compute $\boldsymbol{\lambda}^*$ and $\boldsymbol{\mu}^*$ and set $\mathrm{Cap}(\beta) = \infty$ and $\mathrm{Res}(\beta) = 1$. Therefore, each group can use an unlimited amount of reserved capacity in theory, however each location only has one reserved unit of capacity.

For Algorithm 2, the value of $\Psi(g,t,\beta)$ is set in the following data-driven way: First, we find a confidence interval $[l_g, u_g]$ for the number of remaining arrivals of each group g; this uses the fact that the number of remaining arrivals of each group g follows a Binomial distribution with mean $(T-t)p_g$. Specifically, we set $l_g = N(g,T-t)p_g - Z_g(0.9)$ and $u_g = N(g,T-t)p_g + Z_g(0.9)$ where $Z_g(x)$ denotes the inverse CDF of the group-specific Binomial distribution. Intuitively, for each $n \in [l_g, u_g]$, we want to know: *If there are n remaining arrivals in group g and they are all assigned to their highest-score (greedy) location, what is the chance this group will not meet its minimum requirement?* For each $n \in [l_g, u_g]$, we independently sample n group g arrivals K times. For each of the K length-n vectors, $\{\mathbf{w}_\tau^{k(g)}\}_{\tau=1,\ldots,n}$, we calculate $s_g^k := \sum_{\tau=1}^{n}(\max_j w_{\tau j}^{k(g)} - \mathbb{E}[O_g])$ for $k = 1, \ldots, K$. Finally, we calculate the 10% quantile of $[s_g^1, \ldots, s_g^K]$ and use this value, $q(n,g)$, to set $\Psi(g,t,\beta) = \min_{n\in[l_g,u_g]} q(n,g)$. This effectively lower bounds $\sum_{\tau\in\mathcal{A}(g,T)\setminus\mathcal{A}(g,t)}(\max_j w_{\tau j} - \mathbb{E}[O_g])$ as desired.

BP. To determine the parameters $\boldsymbol{\mu}^*$ used in BP, we compute a sample average of the optimal dual values of the continuous relaxation of OPT. This is also equivalent to the approach described above for ABP without a fairness constraint (*e.g.* achieved by setting $O_g = 0$ for all groups g).

A.2 Additional figures for Section 5.2

Figure 5 shows $\mathrm{UFR}_g(\boldsymbol{\omega})$ for each group under OPT using our three fairness rules. This can be compared to Figure 3 which shows $\mathrm{UFR}_g(\boldsymbol{\omega})$ for RAND and BP.

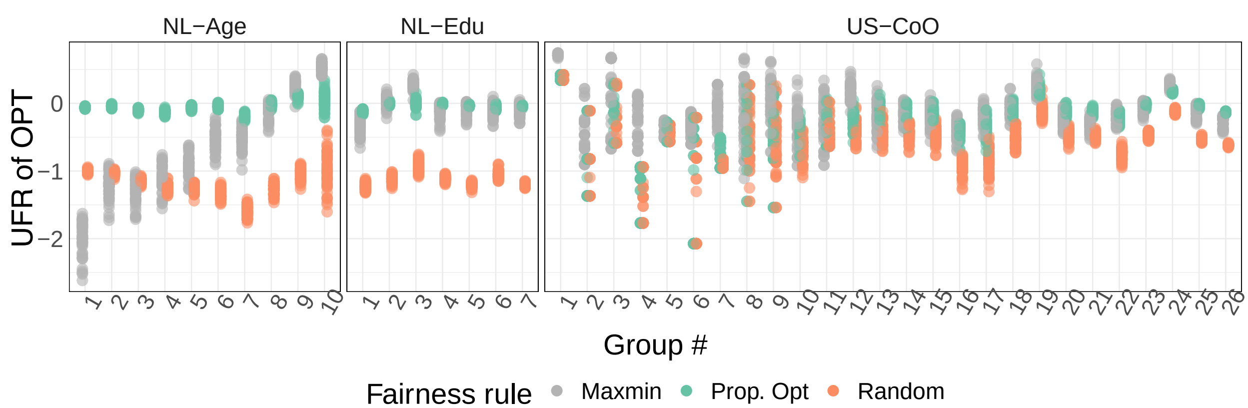

Figure 5: $\mathrm{UFR}_g(\boldsymbol{\omega})$ achieved by OPT with respect to the three fairness targets for each scenario.

A.3 Additional figures and tables for Section 5.3

Figure 6 shows group-level results for the ABP and CBP algorithms, analogous to Figure 3 that was shown for Rand and BP. Figure 7 replicates Figure 4 but aggregates results across the sample paths by taking the average of the median unfairness fairness ratio (conditional on being positive) across groups, instead of the max as in Figure 4. Values of zero in 7 imply that the UFR was non-positive in all sample paths.

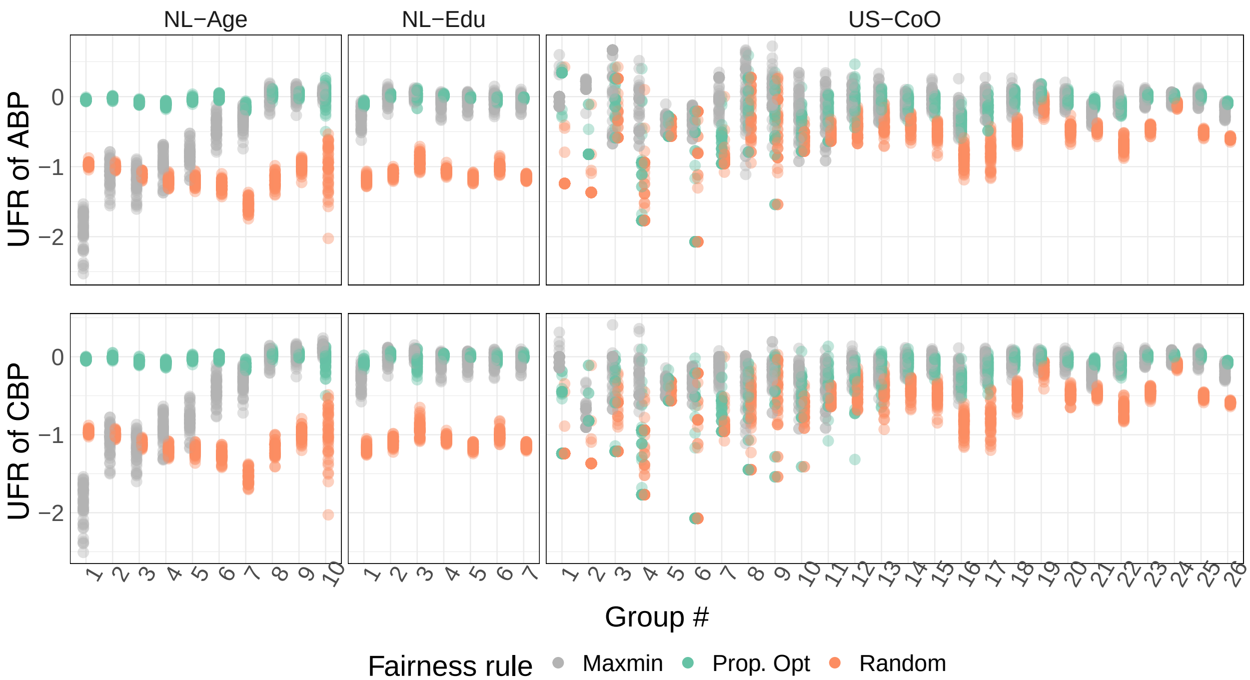

Figure 6: $\mathrm{UFR}_g(\boldsymbol{\omega})$ achieved by ABP and CBP with under our three fairness rules.

A.4 Evaluation beyond i.i.d. based on a single historical trace (Section 5.3)

In this section we demonstrate the performance of our algorithms in a setting that most closely resembles the real world. We treat the 2016 arrivals—in the order that they actually arrived—as the test cohort. Therefore, the arrival sequence no longer satisfies an i.i.d. assumption. All parameters used in the algorithms (e.g., $\boldsymbol{\lambda}^*$, $\boldsymbol{\mu}^*$, and $\mathbb{E}[O_g]$) are determined using the 2015 cohort of arrivals. When running the algorithms, the only knowledge that the algorithm has about the 2016 arrivals is the total number of 2016 arrivals and the capacity vector for 2016. These assumptions reflect the information known in reality.

There is one subtle complexity. Suppose there are certain groups which are present in the 2016 arrival cohort that were not present in the 2015 arrival cohort. This means that we cannot estimate λ_g^* for a group g in this situation. This is not an issue in the NL-Edu or NL-Age scenarios, but is an issue in the US-CoO scenario. We propose handling the issue in the following way. If there is a large group g that arrives in 2016 and not in 2015, it is likely that the policymaker will have prior knowledge about this group. For example, there may have been a policy that prohibited this group from arriving in 2015, but policy changes allow for their arrival in 2016. With this prior knowledge, we could use data-driven approaches to estimate λ_g^* for this group. For example, we

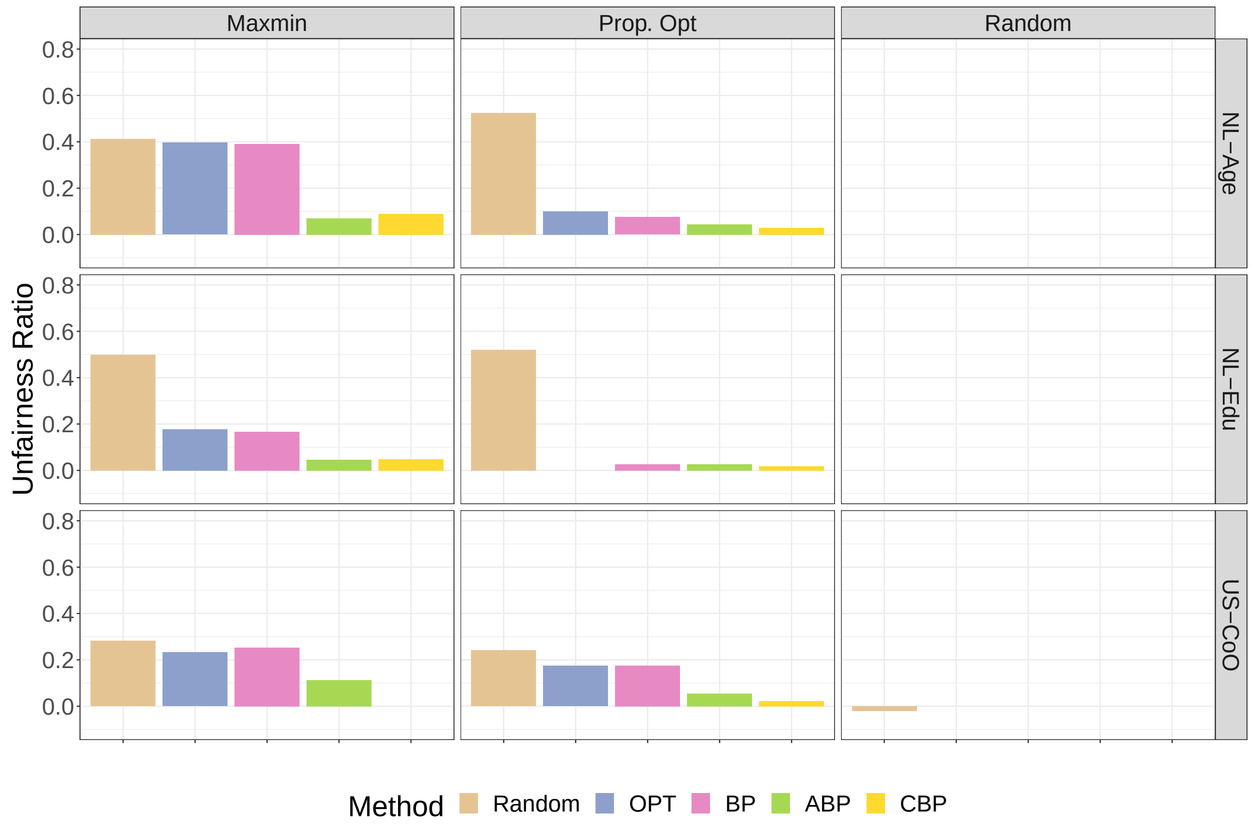

Figure 7: Unfairness ratio aggregated by median. Specifically, for each sample path we take the median unfairness ratio across all groups, and then take the average of the median (conditioned on being positive) across the sample paths. Missing bars indicate that all sample paths had a non-positive UFR.

could generate employment probability predictions for simulated arrivals from this group and add them to the 2015 arrival cohort before computing the parameters of the algorithms.

Now consider a second situation in which refugees from a group g arrive in 2016, but had no arrivals in 2015, and the policymakers did not have prior knowledge about this change. In this case, it is likely that the group is quite small. For example, if a particular group has, on average, only a handful of arrivals, it is likely there are certain years where the group will have zero arrivals, and policymakers might not pay special attention to this. In this situation, we propose either a) using older data that includes arrivals from this group when determining the parameters of the algorithm or b) assigning arrivals to this group to their greedy locations. Because these groups are likely to be quite small, this will not greatly impact the performance of the algorithm. For the numerical results of this section, we take this approach, and assign any arrivals from a group that did not appear in 2015 to their greedy location (among those with remaining capacity).

Group	Size	% positive regret BP	ABP	CBP	Avg unfairness ratio BP	ABP	CBP
1	small	1.00	0.91	0.00	0.41	0.32	0.00
2	small	0.00	0.00	0.00	0.00	0.00	0.00
3	small	0.41	0.39	0.02	0.10	0.09	0.00
4	small	0.02	0.02	0.02	0.00	0.01	0.00
5	small	0.00	0.00	0.00	0.00	0.00	0.00
6	small	0.00	0.00	0.00	0.00	0.00	0.00
7	small	0.00	0.00	0.00	0.00	0.00	0.00
8	small	0.41	0.43	0.00	0.08	0.09	0.00
9	small	0.54	0.35	0.06	0.10	0.05	0.00
10	small	0.00	0.02	0.02	0.00	0.00	0.00
11	small	0.00	0.15	0.04	0.00	0.00	0.00
12	small	0.00	0.28	0.02	0.00	0.03	0.00
13	small	0.26	0.28	0.10	0.01	0.02	0.00
14	medium	0.14	0.12	0.18	0.00	0.00	0.01
15	medium	0.00	0.14	0.02	0.00	0.00	0.00
16	medium	0.00	0.00	0.00	0.00	0.00	0.00
17	medium	0.00	0.10	0.04	0.00	0.00	0.00
18	medium	0.04	0.48	0.16	0.00	0.03	0.01
19	medium	1.00	0.54	0.34	0.17	0.04	0.01
20	medium	0.02	0.36	0.18	0.00	0.01	0.00
21	medium	0.00	0.00	0.00	0.00	0.00	0.00
22	medium	0.00	0.00	0.06	0.00	0.00	0.00
23	large	0.34	0.66	0.74	0.01	0.02	0.01
24	large	1.00	0.56	0.82	0.18	0.01	0.02
25	large	0.06	0.74	0.78	0.00	0.01	0.02
26	large	0.00	0.00	0.00	0.00	0.00	0.00

Table 3: Percentage of instances (out of 50) with positive unfairness ratio for each group, and average unfairness ratio conditional on being positive. "Small" groups are those with fewer than 10 arrivals. "Medium" groups are those with between 10-100 arrivals, and "large" groups have over 100 arrivals.

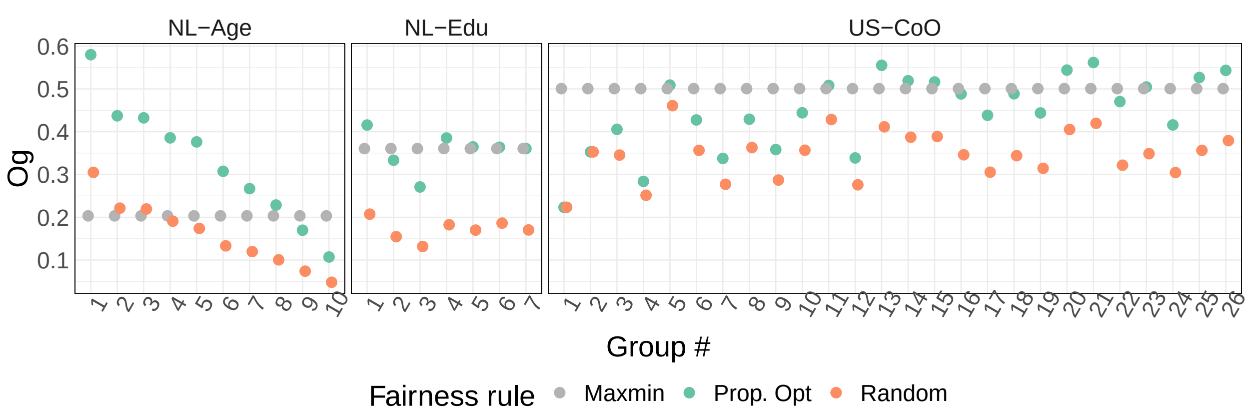

Figure 8: Minimum requirements for the 2016 arrival cohort.

	NL-Edu			NL-Age			US-CoO		
Rule:	Maxmin	PrOpt	Random	Maxmin	PrOpt	Random	Maxmin	PrOpt	Random
RAND	0.46	0.46	0.46	0.46	0.46	0.46	0.67	0.67	0.67
CBP	0.94	0.95	0.96	0.93	0.94	0.96	0.89	0.91	0.92
ABP	0.94	0.96	0.96	0.96	0.96	0.96	0.91	0.92	0.92
BP	0.96	0.96	0.96	0.96	0.96	0.96	0.92	0.92	0.92
OFFLINE$_{\mathcal{F}}$	0.98	1.00	1.00	0.99	1.00	1.00	0.98	1.00	1.00

Table 4: Efficiency of methods in the "real-world" scenario.

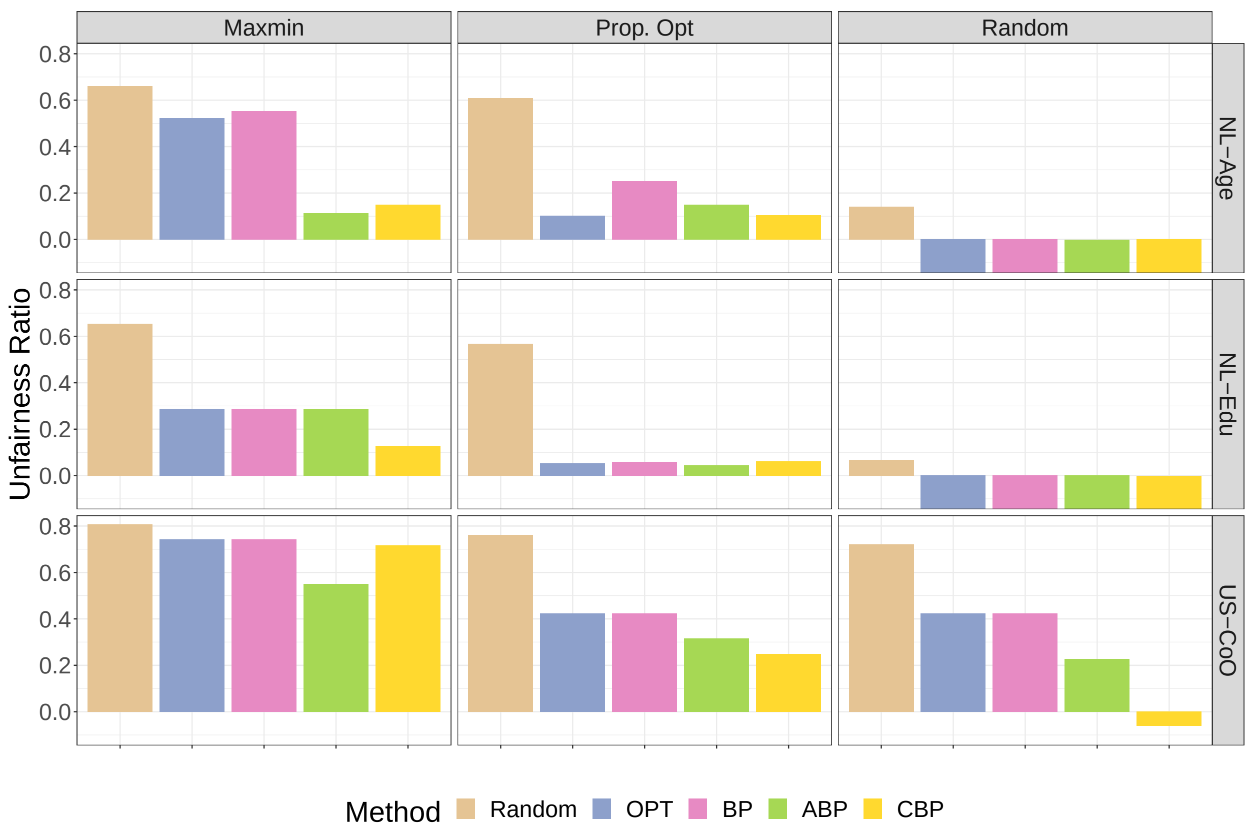

Figure 9: Unfairness ratio (aggregated by max) obtained in the "real-world" scenario.

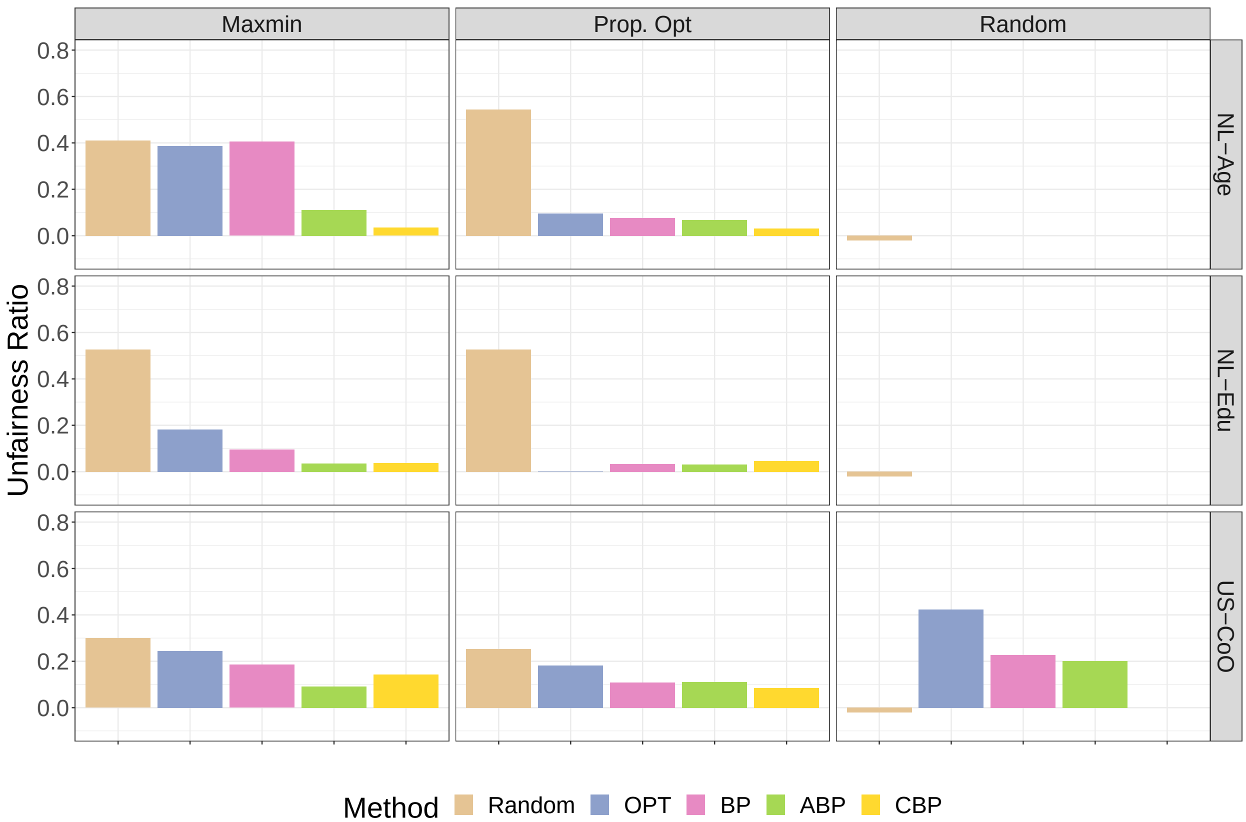

Figure 10: Unfairness ratio (aggregated by median) obtained in the “real-world” scenario.

B Supplementary material for Section 2

B.1 Impossibility sample-path fairness result (Section 2.3)

Let $\Phi(x)$ be the standard normal distribution. We need the following statement of Berry-Esseen Theorem adopted from Theorem 3.4.17 of [Dur19].

Fact 1.

The following proposition shows that there exists a problem setting and a fairness rule such that closely achieving required score or the optimal score on almost every sample path is difficult.

Proposition 1. *There exists a constant $C > 0$, for any $T = 100K$ with sufficiently large integer K, we can construct a problem setting and a fairness rule $\mathcal{F}$ with the property that for any non-anticipatory algorithm, with probability at least C, we must have*

1. $\max_{g\in\mathcal{G}} \left(O_{g,\mathcal{F}}(\boldsymbol{\omega}) - \alpha_g(\boldsymbol{\omega})\right) \geq 0.02$

2. $\max\left(\max_{g\in\mathcal{G}} \mathcal{R}_{g,\mathcal{F}}(\boldsymbol{\omega}), O^\star_{\mathcal{F}}(\boldsymbol{\omega}) - \frac{1}{T}\sum_{t=1}^{T}\sum_{j\in\mathcal{M}} w_{t,j}z_{t,j}\right) \geq 0.02.$

Proof. Fix $T = 100K$ for $K \geq \frac{12}{\Phi(-25)^2}$ and let $C = \frac{(1-4e^{-4})\Phi(-25)}{2} > 0$. Consider a setting with $M = 2$ locations each with capacity $\frac{T}{2}$. There are two groups, each with equal arrival probabilities $p_1 = p_2 = 0.5$. For every case t, if this case is of group 1, then the employment scores are given by $w_{t,1} = 0.57, w_{t,2} = 0$; otherwise, the scores are given by $w_{t,1} = 1, w_{t,2} = 0$. The fairness rule $\mathcal{F}$ is defined as follows. For a sample path $\boldsymbol{\omega}$, the DM observes the number of arrivals of each group and assigns as many arrivals as possible of the larger group to the good location to set the required score. That is, if group 1 has more or equal arrivals than group 2, then $O_{1,\mathcal{F}}(\boldsymbol{\omega}) = \frac{0.285T}{N(1,T)}, O_{2,\mathcal{F}}(\boldsymbol{\omega}) = 0$; otherwise, $O_{1,\mathcal{F}}(\boldsymbol{\omega}) = 0, O_{2,\mathcal{F}}(\boldsymbol{\omega}) = \frac{T}{2N(2,T)}$. Note that we have $\mathbb{E}_{\boldsymbol{\omega}'}[O_{1,\mathcal{F}}(\boldsymbol{\omega}')] \geq \mathbb{E}\left[\frac{0.285T}{T}\mathbb{1}\left(N(1,T) \geq \frac{T}{2}\right)\right] = 0.285\mathbb{P}\{N(1,T) \geq \frac{T}{2}\} = \frac{0.57}{4}$.

Fix a non-anticipatory algorithm. The algorithm could be randomized whose randomness is independent of the case arrival process. We assume the probability space is extended to include the randomness of the algorithm. Define $\tilde{R}(\boldsymbol{\omega})$ to be $\max\left(O_{1,\mathcal{F}}(\boldsymbol{\omega}) - \alpha_1(\boldsymbol{\omega}), O_{2,\mathcal{F}}(\boldsymbol{\omega}) - \alpha_2(\boldsymbol{\omega})\right)$ and let $\hat{R}(\boldsymbol{\omega}) = \max\left(\max_{g\in\mathcal{G}} \mathcal{R}_{g,\mathcal{F}}(\boldsymbol{\omega}), O^\star_{\mathcal{F}}(\boldsymbol{\omega}) - \frac{1}{T}\sum_{t=1}^{T}\sum_{j\in\mathcal{M}} w_{t,j}z_{t,j}\right)$. We next show that with probability at least C, we have $\tilde{R}(\boldsymbol{\omega}), \hat{R}(\boldsymbol{\omega})$ to be greater than 0.02.

Recall that $N(g,t)$ is the number of arrivals of group g up to round t. Let $Q = 0.97T = 97K$. By Hoeffding's Inequality (Fact 3), we have $\mathbb{P}\{|N(1,Q) - \frac{Q}{2}| > 20\sqrt{K}\} \leq 2\exp\left(\frac{-400K}{97K}\right) \leq 2e^{-4}$. Similarly, we have $\mathbb{P}\{|N(2,Q) - \frac{Q}{2}| > 20\sqrt{K}\} \leq 2e^{-4}$. Define

$$\mathcal{S} = \left\{|N(1,Q) - \frac{Q}{2}| \leq 20\sqrt{K}\right\} \cap \left\{|N(2,Q) - \frac{Q}{2}| \leq 20\sqrt{K}\right\}. \tag{3}$$

Then by union bound we have $\mathbb{P}\{\mathcal{S}\} \geq 1 - 4e^{-4}$. In addition, under $\mathcal{S}$ we have $|N(1,Q) - N(2,Q)| \leq 40\sqrt{K}$. Define X by the difference between numbers of arrivals of the two groups in rounds $Q+1$ to T. That is, $X = N(1,T) - N(1,Q) - (N(2,T) - N(2,Q))$. We know X is the sum of $T - Q = 3K$ Rademacher random variables. Then by Berry-Esseen Theorem (Fact 1), we have $|\mathbb{P}\{\frac{X}{\sqrt{3K}} \le a\} - \Phi(a)| \le \sqrt{\frac{3}{K}}$ for any real value a. Taking $a = -25$ leads to $\mathbb{P}\{X \le -41\sqrt{K}\} \ge \Phi(-25) - \sqrt{\frac{3}{K}} \ge \frac{1}{2}\Phi(-25)$ since $K \ge \frac{12}{\Phi(-25)^2}$. By the symmetry of X, we have $\mathbb{P}\{X \ge 41\sqrt{K}\} = \mathbb{P}\{X \le -41\sqrt{K}\} \ge \frac{1}{2}\Phi(-25)$.

Now define A_1 by the number of cases in group 1 that are assigned to location 1 by round Q. Note that both A_1 and $\mathcal{S}$ are independent of X since they are in the filtration of rounds up to Q and are independent of arrivals in rounds $Q+1$ to T. Condition on $\mathcal{S}$ defined in (3). Consider two scenarios:

- $A_1 \ge 0.07T, X \le -41\sqrt{K}$. In this scenario, we know there are more arrivals of group 2 since $N(2,T) - N(1,T) \ge 41\sqrt{K} - 40\sqrt{K} > 0$. By the definition of the fairness rule $\mathcal{F}$, we have $O_{2,\mathcal{F}}(\boldsymbol{\omega})N(2,T) = \frac{T}{2}$. However, since $A_1 \ge 0.07T$, at least $0.07T$ capacity from location 1 are assigned to group 1. We must have $\alpha_2(\boldsymbol{\omega})N(2,T) \le \frac{T}{2} - 0.07T = 43K$, so $\tilde{R}(\boldsymbol{\omega}) \ge \frac{1}{N(2,T)}\left(O_2(\boldsymbol{\omega})N(2,T) - \alpha_2(\boldsymbol{\omega})N(2,T)\right) \ge \frac{7K}{T} = 0.07$. In addition, we know that $O^\star_{\mathcal{F}}(\omega) = \frac{1}{2}$. But since $A_1 \ge 0.07T$, the average global score ALG can obtain is at most $\frac{1}{T}\left(0.57 \times 0.07T + (0.5T - 0.07T)\right) = 0.4699$. As a result, $\hat{\mathcal{R}}(\omega) \ge 0.5 - 0.4699 > 0.03$.
- $A_1 < 0.07T, X \ge 41\sqrt{K}$. Similar to the above scenario, we have that $N(1,T) - N(2,T) > 0$ and thus $O_{1,\mathcal{F}}(\boldsymbol{\omega})N(1,T) = 0.28T$. But since only A_1 cases from group 1 are assigned to location 1 up to round Q and there are at most $T - Q = 3K$ cases after round Q, the total score of group 1 is at most $0.57(0.07T + 0.03T) = 0.057T$. Therefore, under this scenario, $\tilde{R}(\boldsymbol{\omega}) \ge \frac{1}{N(1,T)}\left(O_{1,\mathcal{F}}(\boldsymbol{\omega})N(1,T) - \alpha_1(\boldsymbol{\omega})N(1,T)\right) \ge \frac{1}{T}\left(0.285T - 0.057T\right) = 0.228$. In addition, we know that $\mathcal{R}_{1,\mathcal{F}}(\boldsymbol{\omega}) = \min\left(\mathbb{E}_{\boldsymbol{\omega}'}[O_{1,\mathcal{F}}(\boldsymbol{\omega}')], O_{1,\mathcal{F}}(\boldsymbol{\omega})\right) - \alpha_1(\boldsymbol{\omega}) \ge 0.14 - \alpha_1(\boldsymbol{\omega}) \ge \frac{0.57}{4} - \frac{0.057T}{T/2} \ge 0.02$. As a result, $\hat{\mathcal{R}}(\boldsymbol{\omega}) \ge 0.02$.

Summarizing the above two scenarios, we obtain that

$$
\begin{aligned}
&\mathbb{P}\left\{\min\left(\tilde{\mathcal{R}}(\boldsymbol{\omega}), \hat{\mathcal{R}}(\boldsymbol{\omega})\right) \ge 0.02\right\} \\
&\ge \mathbb{P}\{\mathcal{S}\}\mathbb{P}\left\{\min\left(\tilde{\mathcal{R}}(\boldsymbol{\omega}), \hat{\mathcal{R}}(\boldsymbol{\omega})\right) \ge 0.02 \,\middle|\, \mathcal{S}\right\} \\
&\overset{(a)}{\ge} \mathbb{P}\{\mathcal{S}\}\left(\mathbb{P}\left\{A_1 \ge 0.07T, X \le -41\sqrt{K} \,\middle|\, \mathcal{S}\right\} + \mathbb{P}\left\{A_1 < 0.07T, X \ge 41\sqrt{K} \,\middle|\, \mathcal{S}\right\}\right) \\
&\overset{(b)}{=} \mathbb{P}\{\mathcal{S}\}\left(\mathbb{P}\left\{A_1 \ge 0.07T \,|\, \mathcal{S}\right\}\mathbb{P}\left\{X \le -41\sqrt{K}\right\} + \mathbb{P}\left\{A_1 < 0.07T \,|\, \mathcal{S}\right\}\mathbb{P}\left\{X \ge 41\sqrt{K}\right\}\right) \\
&\overset{(c)}{=} \mathbb{P}\{\mathcal{S}\}\frac{\Phi(-25)}{2}\left(\mathbb{P}\left\{A_1 \ge 0.07T \,|\, \mathcal{S}\right\} + \mathbb{P}\left\{A_1 < 0.07T \,|\, \mathcal{S}\right\}\right) \\
&\overset{(d)}{=} \mathbb{P}\{\mathcal{S}\}\frac{\Phi(-25)}{2} \ge \frac{(1-4e^{-4})\Phi(-25)}{2}
\end{aligned}
$$

where (a) is by the law of total probability and the fact that the two events have no intersection; (b) is by the independence between X and the filtration up to round Q; (c) is by the result we obtained before that $|X| \ge 41\sqrt{K}$ happens with probability at least $\frac{\Phi(-25)}{2}$; (d) is because the two events are complement. We thus finish the proof. $\square$

B.2 Measurability of fairness rule examples (Section 2.2)

For a fairness rule $\mathcal{F}$, our result requires verification of measurability of both the requirement $O_{g,\mathcal{F}}(\boldsymbol{\omega})$ and the associated optimal score $O^\star_{\mathcal{F}}(\boldsymbol{\omega})$. In the sample space, a sample path $\boldsymbol{\omega}$ consists of a discrete sequence $\boldsymbol{gr}(\boldsymbol{\omega}) = (g(t,\boldsymbol{\omega})) \in [G]^T$ and a real vector sequence $\boldsymbol{w}(\boldsymbol{\omega}) = (w_{t,j}(\boldsymbol{\omega})) \in [0,1]^{T\times M}$. Since there are finitely many group arrival sequences, establishing measurability only requires verifying measurability of $O_{g,\mathcal{F}}(\boldsymbol{\omega})$ and $O^\star_{\mathcal{F}}(\boldsymbol{\omega})$ when the group arrival sequence $\boldsymbol{gr}(\boldsymbol{\omega})$ is fixed. As a result, let us consider the sample space $\Omega_{\boldsymbol{gr}}$ where the group arrival sequence $\boldsymbol{gr}(\boldsymbol{\omega})$ is fixed to $\boldsymbol{gr}$. In this scenario, both the required score for a fixed group g, $O_{g,\mathcal{F}}(\boldsymbol{\omega})$, and the optimal score $O^\star_{\mathcal{F}}(\boldsymbol{\omega})$ become functions of the score sequence $\boldsymbol{w}$ and thus are functions over the $T\times M$ Euclidean space. For ease of notations, let $O^{\boldsymbol{gr}}_{g,\mathcal{F}}(\boldsymbol{w}) = O_{g,\mathcal{F}}((\boldsymbol{gr},\boldsymbol{w}))$, $O^{\star;\boldsymbol{gr}}_{\mathcal{F}}(\boldsymbol{w}) = O^\star_{\mathcal{F}}((\boldsymbol{gr},\boldsymbol{w}))$ and denote $\mathcal{D} = (0,1)^{T\times M}$.

It remains to show that $O^{\boldsymbol{gr}}_{g,\mathcal{F}}(\boldsymbol{w})$ and $O^{\star;\boldsymbol{gr}}_{\mathcal{F}}(\boldsymbol{w})$ are Lebesgue measurable functions over $\mathcal{D}$ for the fairness rules under consideration. Note that although the original domain is $[0,1]^{T\times M}$ but not $\mathcal{D}$, we are fine to only consider measurability over $\mathcal{D}$ since $[0,1]^{T\times M}\setminus\mathcal{D}$ is of measure 0. To establish the measurability, we utilize the following result which is a direct application of Theorem 1.1 in [Mar75].

Fact 2. *For $\boldsymbol{w}\in\mathcal{D}$, consider the linear program $M(\boldsymbol{w}) = \max_{\boldsymbol{x}\in\mathbb{R}^{T\times M}} \boldsymbol{c}'(\boldsymbol{w})\boldsymbol{x}$ subject to $\boldsymbol{A}(\boldsymbol{w})\boldsymbol{x} \le \boldsymbol{b}(\boldsymbol{w})$ and $\boldsymbol{x}\ge 0$. Assume that for every $\boldsymbol{w}\in\mathcal{D}$, 1) $\boldsymbol{c}(\boldsymbol{w}),\boldsymbol{A}(\boldsymbol{w}),\boldsymbol{b}(\boldsymbol{w})$ are continuous; 2) $M(\boldsymbol{w})$ exists and is finite 3) the set of optimal solutions is bounded. Then we have $M(\boldsymbol{w})$ is upper semicontinuous at every $\boldsymbol{w}$. If the set of optimal solutions for the dual is also bounded, then $M(\boldsymbol{w})$ is continuous.*

Based on Fact 2, the following lemma shows that $O^{\star;\boldsymbol{gr}}_{\mathcal{F}}(\boldsymbol{w})$ is measurable as long as $O^{\boldsymbol{gr}}_{g,\mathcal{F}}(\boldsymbol{w})$ is continuous over $\mathcal{D}$.

Lemma B.1. *For any group arrival sequence $\boldsymbol{gr}$, suppose $O^{\boldsymbol{gr}}_{g,\mathcal{F}}(\boldsymbol{w})$ is continuous for every group $g\in\mathcal{G}$ and induces feasible* (OFFLINE$_{\mathcal{F}}$) *over $\boldsymbol{w}\in\mathcal{D}$. Then $O^{\star;\boldsymbol{gr}}_{\mathcal{F}}(\boldsymbol{w})$ is measurable over $\mathcal{D}$.*

Proof. Fix $\boldsymbol{w}\in\mathcal{D}$. Recall that $O^{\star;\boldsymbol{gr}}_{\mathcal{F}}(\boldsymbol{w})$ is defined by (OFFLINE$_{\mathcal{F}}$). In the form of Fact 2, it is equivalent to view $\boldsymbol{c}(\boldsymbol{w}) = \boldsymbol{w}$ as a vector and $\boldsymbol{x}$ as the vectorized assignment decision variables $\boldsymbol{z}$ in (OFFLINE$_{\mathcal{F}}$). Let n_g be the number of cases for group g under the sequence $\boldsymbol{gr}$. For constraints, $\boldsymbol{b}(\boldsymbol{w})$ is a $(M+G+T)$-dimensional vector with $\boldsymbol{b}(\boldsymbol{w})_j = s_j$ for $j\in\mathcal{M}$, $\boldsymbol{b}(\boldsymbol{w})_g = -n_g O^{\boldsymbol{gr}}_{g,\mathcal{F}}(\boldsymbol{w})$ for $g\in\mathcal{G}$ and $\boldsymbol{b}(\boldsymbol{w})_t = 1$ for $t\in[T]$. Accordingly, (OFFLINE$_{\mathcal{F}}$) naturally induces the coefficient matrix $\boldsymbol{A}(\boldsymbol{w})$ which is linear in $\boldsymbol{w}$. The objective function $M(\boldsymbol{w})$ is exactly $TO^{\star;\boldsymbol{gr}}_{\mathcal{F}}(\boldsymbol{w})$. We next verify the conditions in Fact 2. First, $\boldsymbol{c}(\boldsymbol{w}),\boldsymbol{A}(\boldsymbol{w})$ are both linear in $\boldsymbol{w}$ and thus continuous. In addition, $\boldsymbol{b}(\boldsymbol{w})$ is continuous since by assumption $O^{\boldsymbol{gr}}_{g,\mathcal{F}}(\boldsymbol{w})$ is continuous for every $g\in\mathcal{G}$. Second, by assumption (OFFLINE$_{\mathcal{F}}$) is feasible and thus $M(\boldsymbol{w}) = TO^{\star;\boldsymbol{gr}}_{\mathcal{F}}(\boldsymbol{w})$ exists. Since the total score is non-negative and at most T, it is also finite. Finally, the set of optimal solutions is bounded since the set of feasible solutions is bounded. Therefore, we have $M(\boldsymbol{w})$ and also $O^{\star;\boldsymbol{gr}}_{\mathcal{F}}(\boldsymbol{w})$ are upper semicontinuous at $\boldsymbol{w}$ by Fact 2. Now, since $O^{\star;\boldsymbol{gr}}_{\mathcal{F}}(\boldsymbol{w})$ is upper semicontinuous at every point of $\mathcal{D}$, it is also measurable over $\mathcal{D}$. □

It remains to show that $O^{\boldsymbol{gr}}_{g,\mathrm{random}}(\boldsymbol{w}), O^{\boldsymbol{gr}}_{g,\mathrm{pro}}(\boldsymbol{w}), O^{\boldsymbol{gr}}_{g,\mathrm{maxmin}}(\boldsymbol{w})$ are all measurable. We show this in the following lemma.

Lemma B.2. *Fix a group arrival sequence $\boldsymbol{gr}$ and a group g. We have that $O^{\boldsymbol{gr}}_{g,\mathrm{random}}(\boldsymbol{w}), O^{\boldsymbol{gr}}_{g,\mathrm{pro}}(\boldsymbol{w})$, $O^{\boldsymbol{gr}}_{g,\mathrm{maxmin}}(\boldsymbol{w})$ are continuous over $\mathcal{D}$.*

Proof. By definition, $O^{\boldsymbol{gr}}_{g,\mathrm{random}}(\boldsymbol{w}) = \frac{1}{|\mathcal{A}(g,T)|}\sum_{t\in\mathcal{A}(g,T)}\sum_{j\in\mathcal{M}}\frac{s_j w_{t,j}}{\sum_{j'\in\mathcal{M}} s_{j'}}$ where $\mathcal{A}(g,T)$ is a fixed set since $\boldsymbol{gr}$ is fixed. Therefore, $O^{\boldsymbol{gr}}_{g,\mathrm{random}}(\boldsymbol{w})$ is a linear function of $\boldsymbol{w}$ and thus is continuous.

Recall that $\mathcal{Z}_g = \{\boldsymbol{z} \in [0,1]^{\mathcal{A}(g,T)\times M} : \sum_{j\in\mathcal{M}} z_{t,j} \le 1, \forall t \in \mathcal{A}(g,T); \; \sum_{t\in\mathcal{A}(g,T)} z_{t,j} \le s_{j,g}, \forall j \in \mathcal{M}\}$ for $O^{\boldsymbol{gr}}_{g,\text{pro}}(\boldsymbol{w})$. When the group arrival sequence is fixed, $\mathcal{Z}_g$ is a fixed polytope. Let $\mathcal{E}_g$ be the finite set of extreme points of $\mathcal{Z}_g$. We have $O^{\boldsymbol{gr}}_{g,\text{pro}}(\boldsymbol{w}) = \frac{1}{N(g,T)}\max_{\boldsymbol{z}\in\mathcal{Z}_g} \sum_{t\in\mathcal{A}(g,T)} \sum_{j\in\mathcal{M}} w_{t,j} z_{t,j} = \frac{1}{N(g,T)}\max_{\boldsymbol{z}\in\mathcal{E}} \sum_{t\in\mathcal{A}(g,T)} \sum_{j\in\mathcal{M}} w_{t,j} z_{t,j}$. Since the last one is the maximum of finitely many linear function of $\boldsymbol{w}$, we conclude that $O^{\boldsymbol{gr}}_{g,\text{pro}}(\boldsymbol{w})$ is continuous.

Finally, let us consider $O^{\boldsymbol{gr}}_{g,\text{maxmin}}(\boldsymbol{w})$. For ease of notations, define $\mathcal{G}'$ as the set of non-empty groups in the fixed group arrival sequence $\boldsymbol{gr}$. Then by definition, $O^{\boldsymbol{gr}}_{g,\text{maxmin}}(\boldsymbol{w})$ is given by $\max_{\boldsymbol{z}\in\mathcal{Z}} \min_{g'\in\mathcal{G}'} \frac{1}{N(g',T)} \sum_{t\in\mathcal{A}(g',T)} \sum_{j=1}^{M} w_{t,j} z_{t,j}$. It can be equivalently stated by the following program:

$$
\begin{aligned}
&\max_{\boldsymbol{z}\in\mathbb{R}_+^{T\times M},\varphi\in\mathbb{R}_+} \quad \varphi \\
&s.t. \quad \varphi \le \frac{1}{N(g',T)} \sum_{t\in\mathcal{A}(g',T)} \sum_{j\in\mathcal{M}} w_{t,j} z_{t,j}, \quad \forall g' \in \mathcal{G}' \\
&\qquad \sum_{t\in[T]} z_{t,j} \le s_j, \quad \forall j \in \mathcal{M}, \qquad \sum_{j=1}^{M} z_{t,j} \le 1, \quad \forall t \in [T]
\end{aligned}
\tag{4}
$$

The first three conditions in Fact 2 hold naturally. To show that $O^{\boldsymbol{gr}}_{g,\text{maxmin}}(\boldsymbol{w})$ is continuous, it remains to prove that the dual of (4) has a bounded set of optimal solutions for a fixed $\boldsymbol{w}$. Let $\{u_g\}, \{x_j\}, \{y_t\}$ denote the corresponding dual variables for constraints in (4). Its dual is given by

$$
\begin{aligned}
&\min_{\boldsymbol{u}\in\mathcal{R}_+^{|\mathcal{G}'|},\boldsymbol{x}_+^M,\boldsymbol{y}_+^T} \quad \sum_{j\in\mathcal{M}} x_j s_j + \sum_{t\in[T]} y_t \\
&s.t. \quad \sum_{g'\in\mathcal{G}'} u_{g'} \ge 1, \quad -\frac{w_{t,j} u_{g(t)}}{N(g(t),T)} + x_j + y_t \ge 0, \quad \forall t \in [T], j \in \mathcal{M}
\end{aligned}
\tag{5}
$$

Denote $\varphi^\star$ by the optimal value to (4) which exists and is finite. Let $(\boldsymbol{u}^\star, \boldsymbol{x}^\star, \boldsymbol{y}^\star)$ be any optimal solution to the dual (5). Recall that $\boldsymbol{w} \in \mathcal{D}$ and thus $w_{t,j} > 0$ for every $t \in [T], j \in \mathcal{M}$. We next show that $(\boldsymbol{u}^\star, \boldsymbol{x}^\star, \boldsymbol{y}^\star) \in \left[0, \frac{2T\varphi^\star}{\min_{t\in[T],j\in\mathcal{M}} w_{t,j}}\right]^{|\mathcal{G}'|+M+T}$. First, its elements are non-negative by definition. Second, since it is an optimal solution, we have $\sum_{j\in\mathcal{M}} x_j^\star s_j + \sum_{t\in[T]} y_t = \varphi^\star$. As a result, $x_j^\star, y_t^\star \le \varphi^\star$ for any $j \in \mathcal{M}, t \in [T]$. Now take any group $g' \in \mathcal{G}'$. Since it is non-empty, there exists a case t' such that $g(t') = g'$. Take any location $j' \in \mathcal{M}$. By the constraint for τ and location j' in (5), we have $-\frac{w_{t',j'} u^\star_{g'}}{N(g',T)} + x^\star_{j'} + y^\star_{t'} \ge 0$. Since $x^\star_{j'}, y^\star_{t'} \le \varphi$ and $1 \le N(g',T) \le T$, we have $u^\star_{g'} \le \frac{2T\varphi^\star}{w_{t',j'}}$. Extending the result to every group in $\mathcal{G}'$ shows that the set of optimal solutions for the dual program (5) is bounded. Using Fact 2 concludes the proof by proving $O^{\boldsymbol{gr}}_{g,\text{maxmin}}(\boldsymbol{w})$ is continuous over $\mathcal{D}$. $\square$

Combining Lemmas B.1 and B.2 gives the desired measurability result.

Proposition 2. *The Random, Proportionally Optimized, and Maxmin fairness rules give measurable required score $O_{g,\mathcal{F}}$ and optimal score $O^\star_{\mathcal{F}}$.*

Proof. Let $\mathcal{F}$ be any of the three fairness rules. By definition of the rule, (OFFLINE$_{\mathcal{F}}$) is always feasible. Fix a group arrival sequence $\boldsymbol{gr}$. Lemma B.2 shows that $O^{\boldsymbol{gr}}_{g,\mathcal{F}}$ is continuous over $\mathcal{D}$ for every group $g \in \mathcal{G}$. Lemma B.1 then gives the measurability of $O^{\star,\boldsymbol{gr}}$ over $\mathcal{D}$. Now since there is only finitely many group arrival sequences (from $[G]^T$), we have $O^\star_{\mathcal{F}}(\boldsymbol{\omega}), \{O_{g,\mathcal{F}}(\boldsymbol{\omega})\}$ are measurable over $[G]^T \times \mathcal{D}$. We then finish the proof by noting that they are also measurable over the original sample space (with Borel $\sigma-$algebra) since $[0,1]^{T\times M} \setminus \mathcal{D}$ is of measure 0. $\square$

C Supplementary material for Section 3

C.1 Guarantee of global regret for Amplified Bid Price Control (Theorem 1)

Proof of Theorem 1. Since the algorithm selects location $J^{\mathrm{WAB}}(t)$ for every case $t \leq T_{\mathrm{emp}}$, we have the average employment score $\frac{1}{T}\sum_{t\in[T]}\sum_{j\in\mathcal{M}} w_{t,j}z_{t,j}$ is lower bounded by $\frac{1}{T}\sum_{t=1}^{T_{\mathrm{emp}}} w_{t,J^{\mathrm{WAB}}(t)} \geq \frac{1}{T}\sum_{t=1}^{T} w_{t,J^{\mathrm{WAB}}(t)} - \frac{T-T_{\mathrm{emp}}}{T}$. Due to the i.i.d. nature of cases, By Hoeffding's Inequality (Fact 3), with probability at least $1-\frac{\delta}{M+2}$, we have $\frac{1}{T}\sum_{t=1}^{T} w_{t,J^{\mathrm{WAB}}(t)} \geq \mathbb{E}\left[\sum_{t=1}^{T} w_{t,J^{\mathrm{WAB}}(t)}\right] - \sqrt{\frac{1}{2}T\ln\left(\frac{M+2}{\delta}\right)}$ which is lower bounded by $T\mathbb{E}\left[O^\star\right] - \sqrt{\frac{1}{2}T\ln\left(\frac{M+2}{\delta}\right)}$ by Lemma 3.1. Lemma 3.3 shows that with probability at least $1-\frac{M\delta}{M+2}$, $T_{\mathrm{emp}} \geq T - \Delta^{\mathrm{WAB}}$. As a result of union bound, with probability at least $1-\delta$, $\frac{1}{T}\sum_{t\in[T]}\sum_{j\in\mathcal{M}} w_{t,j}z_{t,j} \geq \mathbb{E}\left[O^\star\right] - \frac{\Delta^{\mathrm{WAB}}+\sqrt{\frac{1}{2}T\ln\left(\frac{M+2}{\delta}\right)}}{T}$ and thus $\mathcal{R}^{\mathrm{ABP}}_{\mathcal{F}} = \mathbb{E}\left[O^\star\right] - \frac{1}{T}\sum_{t\in[T]}\sum_{j\in\mathcal{M}} w_{t,j}z_{t,j} \leq \frac{\Delta^{\mathrm{WAB}}+\sqrt{\frac{1}{2}T\ln\left(\frac{M+2}{\delta}\right)}}{T}$. Note that if $M=1$, the regret is always zero and thus we can assume $M \geq 2$ under which $\frac{M+2}{\delta} \leq \left(\frac{M}{\delta}\right)^2$ for $\delta < 1$ (the theorem automatically holds when $\delta \geq 1$). As a result, plugging in the value of Δ^{WAB}, we have $\mathcal{R}^{\mathrm{ABP}}_{\mathcal{F}} \leq \sqrt{\frac{\ln(M/\delta)}{T}}\left(\frac{1}{\hat{s}_{\min}}+1\right)$ with probability at least $1-\delta$. $\square$

C.2 Guarantee of g-regret for Amplified Bid Price Control (Theorem 2)

Proof of Theorem 2. For ease of notation, we set $\beta = \left(\frac{\delta}{M+2}\right)^{1/4}$ in the following analysis. Fix a group $g \in \mathcal{G}$. To bound $\mathcal{R}^{\mathrm{ABP}}_{g,\mathcal{F}}$, consider the event $\mathcal{S}_1$ where $T_{\mathrm{emp}} \geq T-\Delta^{\mathrm{WAB}}$. Condition on $\mathcal{S}_1$, we have $|\mathcal{A}(g,T)|\alpha_g \geq \sum_{t=1}^{T-\Delta^{\mathrm{WAB}}} w_{t,J^{\mathrm{WAB}}(t)}\mathbb{1}\left(g(t)=g\right)$. Let $X_t = w_{t,J^{\mathrm{WAB}}(t)}\mathbb{1}\left(g(t)=g\right)$ for $t \in [T]$. Note that $\{X_t\}_{t\in[T]}$ are i.i.d. non-negative random variables with $\mathbb{E}\left[X_t^2\right] \leq p_g$. We can apply a variant of Bernstein's Inequality (see Lemma E.1) to get that with probability at least $1-\beta^4$, it holds

$$\sum_{t=1}^{T-\Delta^{\mathrm{WAB}}} X_t \geq \mathbb{E}\left[\sum_{t=1}^{T-\Delta^{\mathrm{WAB}}} X_t\right] - 2\sqrt{2p_g(T-\Delta^{\mathrm{WAB}})\ln(1/\beta)}.$$

Denote the above event by $\mathcal{S}_2$. Note that

$$\mathbb{E}\left[X_t\right] = \mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)}\mathbb{1}\left(g(t)=g\right)\right] = \mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)}\,\middle|\,g(t)=g\right]\mathbb{P}\{g(t)=g\}.$$

Lemma 3.2 shows that $\mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)}\,\middle|\,g(t)=g\right] \geq \mathbb{E}\left[O_g\right] - \sqrt{\frac{1}{p_gT}}$. Therefore, we have $\mathbb{E}\left[X_t\right] \geq p_g\mathbb{E}\left[O_g\right] - \sqrt{\frac{p_g}{T}}$. As a result of linearity of expectation, under $\mathcal{S}_1 \cap \mathcal{S}_2$, the total score of group g is lower bounded by

$$\begin{aligned}|\mathcal{A}(g,T)|\alpha_g &\geq (T-\Delta^{\mathrm{WAB}})\left(p_g\mathbb{E}\left[O_g\right] - \sqrt{p_g/T}\right) - 2\sqrt{2p_g(T-\Delta^{\mathrm{WAB}})\ln(1/\beta)}\\ &\geq (T-\Delta^{\mathrm{WAB}})p_g\mathbb{E}\left[O_g\right] - 3\sqrt{2p_gT\ln(1/\beta)}\end{aligned}$$

By an implication of Chernoff bound (see Lemma E.2), with probability at least $1-\beta^4$, the number of group g cases is at most $|\mathcal{A}(g,T)| \le p_g T + \xi$ where $\xi = 5\sqrt{\max(p_g T, \ln(1/\beta))\ln(1/\beta)}$. Denote this event by $\mathcal{S}_3$. Then under $\mathcal{S}_1 \cap \mathcal{S}_2 \cap \mathcal{S}_3$, we have

$$\alpha_g \ge \frac{(T-\Delta^{\mathrm{WAB}})p_g \mathbb{E}[O_g] - 3\sqrt{2p_g T\ln(1/\beta)}}{p_g T + \xi} \ge \frac{(T-\Delta^{\mathrm{WAB}})p_g \mathbb{E}[O_g]}{p_g T + \xi} - 3\sqrt{\frac{2\ln(1/\beta)}{p_g T}}.$$

Now note that

$$\frac{(T-\Delta^{\mathrm{WAB}})p_g \mathbb{E}[O_g]}{p_g T + \xi} = \left(1 - \frac{\Delta^{\mathrm{WAB}} p_g + \xi}{p_g T + \xi}\right)\mathbb{E}[O_g] \ge \mathbb{E}[O_g] - \frac{\Delta^{\mathrm{WAB}}}{T} - \frac{\xi}{p_g T}$$

and $\frac{\xi}{p_g T} \le 5\sqrt{\frac{\ln(1/\beta)}{p_g T}} + \frac{5\ln(1/\beta)}{p_g T}$. As a result, under $\mathcal{S}_1 \cap \mathcal{S}_2 \cap \mathcal{S}_3$, we have α_g is lower bounded by

$$\mathbb{E}[O_g] - \frac{\Delta^{\mathrm{WAB}}}{T} - 5\sqrt{\frac{\ln(1/\beta)}{p_g T}} + \frac{5\ln(1/\beta)}{p_g T} - 3\sqrt{\frac{2\ln(1/\beta)}{p_g T}} \ge \mathbb{E}[O_g] - \frac{\Delta^{\mathrm{WAB}}}{T} - 10\sqrt{\frac{\ln(1/\beta)}{p_g T}} - \frac{5\ln(1/\beta)}{p_g T}$$

As a result, $\mathcal{R}^{\mathrm{ABP}}_{g,\mathcal{F}} \le \mathbb{E}[O_g] - \alpha_g \le \sqrt{\frac{2\ln(1/\beta)}{T}}\frac{1}{\hat{s}_{\min}} + 10\sqrt{\frac{\ln(1/\beta)}{p_g T}} + \frac{5\ln(1/\beta)}{p_g T}$. Note that $\mathcal{R}^{\mathrm{ABP}}_{g,\mathcal{F}}$ is always upper bounded by 1. Therefore,

$$\begin{aligned}\mathcal{R}^{\mathrm{ABP}}_{g,\mathcal{F}} &\le \sqrt{\frac{2\ln(1/\beta)}{T}}\frac{1}{\hat{s}_{\min}} + \min\left(10\sqrt{\frac{\ln(1/\beta)}{p_g T}} + \frac{5\ln(1/\beta)}{p_g T}, 1\right) \\ &\le \sqrt{\frac{2\ln(1/\beta)}{T}}\frac{1}{\hat{s}_{\min}} + 20\sqrt{\frac{\ln(1/\beta)}{p_g T}} = \sqrt{\frac{2\ln(1/\beta)}{T}}\left(\frac{1}{\hat{s}_{\min}} + \sqrt{\frac{200}{p_g}}\right)\end{aligned}$$

where the second inequality is due to the fact that if $\frac{5\ln(1/\beta)}{p_g T} > 10\sqrt{\frac{\ln(1/\beta)}{p_g T}}$, we would have $10\sqrt{\frac{\ln(1/\beta)}{p_g T}} \ge 20$ and thus the minimum term is equal to 1. Note that $\mathbb{P}\{\mathcal{S}_1\} \ge 1 - M\beta^4$ by Lemma 3.3 and $\mathbb{P}\{\mathcal{S}_2\}, \mathbb{P}\{\mathcal{S}_3\} \ge 1-\beta^4$. By the union bound, we have $\mathbb{P}\{S_1 \cap S_2 \cap S_3\} \ge 1-(M+2)\beta^4$. As a result, for any $\beta > 0$, with probability at least $1-(M+2)\beta^4$, we have $\mathcal{R}^{\mathrm{ABP}}_{g,\mathcal{F}} \le \sqrt{\frac{2\ln(1/\beta)}{T}}\left(\frac{1}{\hat{s}_{\min}} + \sqrt{\frac{200}{p_g}}\right)$.

Since $\beta = (\delta/(M+2))^{1/4}$ and we can safely assume $M \ge 2, \delta < 1$ (otherwise the theorem automatically holds), we have $\sqrt{\frac{2\ln(1/\beta)}{T}} \le \sqrt{\frac{\ln(M/\delta)}{T}}$ and $1-(M+2)\beta^4 = 1-\delta$. As a result, with probability at least $1-\delta$, $\mathcal{R}^{\mathrm{ABP}}_{g,\mathcal{F}} \le \sqrt{\frac{\ln(M/\delta)}{T}}\left(\frac{1}{\hat{s}_{\min}} + \sqrt{\frac{200}{p_g}}\right)$. □

C.3 Guarantee for ABP when no capacity constraint (Lemmas 3.1, 3.2)

The proofs rely on KKT conditions of solving the convex optimization problem for $\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star$ with similar ideas from [TVR98]. Recall the minimization problem in line 1 in Amplified Bid Price Control that solves $\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star$. We can write out its Lagrangian defined by

$$\begin{aligned}L(\boldsymbol{\mu}, \boldsymbol{\lambda}, \boldsymbol{u}^1, \boldsymbol{u}^2) = T\mathbb{E}_{\boldsymbol{\theta}}\left[\max_{j\in\mathcal{M}}((1+\lambda_{g(\boldsymbol{\theta})})w_j(\boldsymbol{\theta}) - \mu_j)\right] + \sum_{j\in\mathcal{M}}\mu_j s_j - \sum_{g\in\mathcal{G}}\lambda_g\mathbb{E}[O_g N(g,T)] \\ - \sum_{j\in\mathcal{M}}\mu_j u^1_j - \sum_{g\in\mathcal{G}}\lambda_g u^2_g,\end{aligned} \tag{6}$$

where $\boldsymbol{u}^1 \in \mathbb{R}_+^M, \boldsymbol{u}^2 \in \mathbb{R}_+^G$ are associated Lagrange multipliers. Note that $L(\boldsymbol{\mu}, \boldsymbol{\lambda}, \boldsymbol{u}^1, \boldsymbol{u}^2)$ is continuously differentiable because the number of groups is finite and condition on a group, the score distribution is continuous. By the Karush–Kuhn–Tucker (KKT) conditions, a necessary condition for $\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star$ to be optimal is that the partial derivatives, $\frac{\partial L}{\partial \boldsymbol{\mu}}, \frac{\partial L}{\partial \boldsymbol{\lambda}}$ are both zero. We then have for every $j \in \mathcal{M}, g \in \mathcal{G}, t \in T$,

$$-T\mathbb{P}\{J^{\mathrm{WAB}}(t) = j\} + s_j - u_j^1 = 0;\; Tp_g \mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)} \,\middle|\, g(1) = g\right] - \mathbb{E}\left[O_g N(g,T)\right] - u_g^2 = 0. \tag{7}$$

Proof of Lemma 3.2. Since $u_g^2 \geq 0$, we have $\mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)} \,\middle|\, g(t) = g\right] \geq \frac{1}{p_g T}\mathbb{E}\left[O_g N(g,T)\right]$ by (7). Furthermore, note that by Cauchy-Schwarz inequality,

$$\left|\mathbb{E}\left[O_g N(g,T)\right] - \mathbb{E}\left[O_g\right]\mathbb{E}\left[N(g,T)\right]\right| \leq \sqrt{\mathrm{var}(O_g)\mathrm{var}(N(g,T))} \leq \sqrt{\mathrm{var}(N(g,T))} \leq \sqrt{p_g T}.$$

As a result, $\frac{1}{p_g T}\mathbb{E}\left[O_g N(g,T)\right] \geq \mathbb{E}\left[O_g\right] - \sqrt{\frac{1}{p_g T}}$ and thus $\mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)} \,\middle|\, g(t) = g\right] \geq \mathbb{E}\left[O_g\right] - \sqrt{\frac{t}{p_g T}}$. We complete the proof by noting weights of cases are i.i.d. $\square$

We next show $\mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)}\right] \geq \mathbb{E}\left[O^\star\right]$ for Lemma 3.1.

Proof of Lemma 3.1. Recall the Lagrangian relaxation in (LAGR). By definition of $L(\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star)$, we have

$$\begin{aligned}
\mathbb{E}\left[L(\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star)\right] &= \mathbb{E}\left[\sum_{t=1}^{T} \max_{j \in \mathcal{M}}((1 + \lambda_{g(t)}^\star) w_{t,j} - \mu_j^\star) + \sum_{j \in \mathcal{M}} \mu_j^\star s_j - \sum_{g \in \mathcal{G}} \lambda_g^\star O_g N(g,T)\right] \\
&= \sum_{t=1}^{T} \mathbb{E}\left[(1 + \lambda_{g(t)}^\star) w_{t,J^{\mathrm{WAB}}(t)} - \mu_{J^{\mathrm{WAB}}(t)}^\star\right] + \sum_{j \in \mathcal{M}} \mu_j^\star s_j - \sum_{g \in \mathcal{G}} \lambda_g^\star \mathbb{E}\left[O_g N(g,T)\right] \\
&= \sum_{t \in [T]} \mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)}\right] + \sum_{j \in \mathcal{M}} \mu_j^\star (s_j - \sum_{t \in [T]} \mathbb{P}\{j = J^{\mathrm{WAB}}(t)\}) \\
&\quad - \sum_{g \in \mathcal{G}} \lambda_g^\star \left(\mathbb{E}\left[O_g N(g,T)\right] - \sum_{t \in [T]} p_g \mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)} \,\middle|\, g(t) = g\right]\right).
\end{aligned} \tag{8}$$

where the second equality uses linearity of expectation; the third equality uses linearity of expectation again and the i.i.d. assumption of arrivals. Recall that $(\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star)$ is an optimal solution to the constrained optimization of $\min \mathbb{E}\left[L(\boldsymbol{\mu}, \boldsymbol{\lambda})\right]$. We now apply the complementary slackness in KKT conditions on the Lagrangian multipliers (6). For every location $j \in \mathcal{M}$ and group $g \in \mathcal{G}$, we have that $\mu_j^\star u_j^1 = 0, \lambda_g^\star u_g^2 = 0$. Combining this fact with (7), we have that $\forall j \in \mathcal{M}, g \in \mathcal{G}, t \in [T]$,

$$\mu_j^\star (s_j - T\mathbb{P}\{j = J^{\mathrm{WAB}}(t)\}) = 0,\; \lambda_g^\star \left(\mathbb{E}\left[O_g N(g,T)\right] - Tp_g \mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(t)} \,\middle|\, g(t) = g\right]\right).$$

Putting it back to (8) gives us that $\sum_{t \in [T]} \mathbb{E}\left[w_{t,J^{\mathrm{WAB}}(T)}\right] = \mathbb{E}\left[L(\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star)\right]$. Since the Lagrangian satisfies $L(\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star) \geq TO^\star$ for all sample paths, we also have $\mathbb{E}\left[L(\boldsymbol{\mu}^\star, \boldsymbol{\lambda}^\star)\right] \geq \mathbb{E}\left[TO^\star\right]$, which finishes the proof by the i.i.d. nature of arrivals. $\square$

C.4 Lower bound T_{emp} under Amplified Bid Price Control (Lemma 3.3)

We first give the following lemma on the probability that a case is sent to a particular location under J^{WAB}. It is proved using the KKT condition.

Lemma C.1. *For every location $j \in \mathcal{M}$ and case $t \in [T]$, we have $\mathbb{P}\{J^{\text{WAB}}(t) = j\}T \leq s_j$.*

Proof of Lemma C.1. Since $u_j^1 \geq 0$, we immediately have $T\mathbb{P}\{J^{\text{WAB}}(t) = j\} \leq s_j$ by (7). The result follows from the i.i.d. nature of cases. □

Proof of Lemma 3.3. Recall that $a_j(t)$ is the consumed capacity in location j in the first t periods. Then $T_{\text{emp}} = \min_t\{\exists j \in \mathcal{M}, a_j(t) = s_j(t)\}$. Fix a location $j \in \mathcal{M}$. Let us set $b_j(t) = 1$ if $J^{\text{WAB}}(t) = j$. Since Amplified Bid Price Control follows J^{WAB} in the first T_{emp} periods, we have for all $t \leq T_{\text{emp}}$, $a_j(t) = \sum_{\tau=1}^t b_j(\tau)$. Notice that $\{b_j(\tau), \tau \in [T]\}$ are i.i.d. Bernoulli random variables. Using Fact 3, for a fixed case t, we have

$$\mathbb{P}\left\{\sum_{\tau=1}^t b_j(\tau) > t\mathbb{P}\{J^{\text{WAB}}(t) = j\} + \sqrt{\frac{1}{2}t\ln\left(\frac{M+2}{\delta}\right)}\right\} \leq \frac{\delta}{M+2}. \tag{9}$$

Consider time $T_1 = T - \hat{s}_{\min}^{-1}\sqrt{\frac{1}{2}T\ln\left(\frac{M+2}{\delta}\right)}$. Applying (9) and the union bound over all $j \in \mathcal{M}$, with probability at least $1 - \frac{M\delta}{M+2}$, for every $j \in \mathcal{M}$, we have

$$\sum_{\tau=1}^{T_1} b_j(\tau) \leq T_1\mathbb{P}\{J^{\text{WAB}}(1) = j\} + \sqrt{\frac{1}{2}T\ln\left(\frac{M+2}{\delta}\right)} \leq T_1\hat{s}_j + \sqrt{\frac{1}{2}T\ln\left(\frac{M+2}{\delta}\right)} \leq s_j,$$

where the second inequality is by Lemma C.1 that $\mathbb{P}\{J^{\text{WAB}}(1) = j\} \leq \frac{s_j}{T} = \hat{s}_j$. As a result, with probability at least $1 - \frac{M}{M+2}\delta$, we have $T_{\text{emp}} \geq T_1 = T - \hat{s}_{\min}^{-1}\sqrt{\frac{1}{2}T\ln\left(\frac{M+2}{\delta}\right)}$. □

D Supplementary material for Section 4

D.1 Guarantee of g-regret for Conservative Bid Price Control (Theorem 4)

For Theorem 4, we set

$$T_1 = \frac{10^9(M+G)^4}{\hat{s}_{\min}^3\varepsilon^3\delta^2}. \tag{10}$$

Recall that $\Delta^{\text{WCB}} = \hat{s}_{\min}^{-1}\left(\frac{11\sqrt{GT}\ln(1/\beta)}{\varepsilon} + \frac{32G\ln(1/\beta)}{\varepsilon^2}\right)$. We first give a result based on fixed conservative parameter β.

Lemma D.1. *Given an ex-post feasible fairness rule $\mathcal{F}$ with slackness ε, suppose that $T \geq 36\Delta^{\text{WCB}}$. Then For each group $g \in \mathcal{G}$, with probability at least $1 - (5T+5)(M+G)\beta^4 - \frac{378MG\ln^3(1/\beta)}{\varepsilon^2 T}$, we have $\mathcal{R}_{g,\mathcal{F}}^{\text{CBP}} \leq \frac{12\ln(1/\beta)}{\hat{s}_{\min}\varepsilon\sqrt{T}}\left(11\sqrt{G} + \frac{32G}{\varepsilon\sqrt{T}}\right)$.*

Proof. Lemma 4.3 shows the regret guarantee for a group g with $p_g \geq \frac{54\ln^2(1/\beta)}{\varepsilon_g^2 T}$. We then finish the proof by using Lemma 4.4 for group g with $p_g < \frac{54\ln^2(1/\beta)}{\varepsilon_g^2 T}$. □

We next provide a lemma concerning the comparison between a logarithm and a polynomial.

Lemma D.2. *For $x \geq 100$, we have $\ln^3(x) \leq x$. For $x \geq 10^9$, we have $\ln(x) \leq x^{1/6}$.*

Proof. We first prove the first result. Let $h(x) = x - \ln^3(x)$ with domain $x \in (0, +\infty)$. Then $h'(x) = 1 - \frac{3\ln^2(x)}{x}$ and $h''(x) = \frac{3(\ln x - 2)\ln x}{x^2}$. We observe that $h''(x) > 0$ for $x > e^2$ and thus $h'(x)$ increases in $(e^2, +\infty)$. Now note that $h'(99) > 0$; as a result, $h'(x) > 0$ for any $x > 99$. Therefore, $h(x)$ increases in $(99, +\infty)$. Since $h(100) > 0$, we know $x \geq \ln^3(x)$ for any $x > 100$.

For the second result, let $q(x) = x^{1/6} - \ln(x)$. Then $q'(x) = \frac{1}{6}x^{-5/6} - \frac{1}{x}$ and $q'(x) > 0$ for any $x > 6^6 = 46656$. Therefore $q(x)$ increases in $(46656, +\infty)$. Since $q(10^9) > 0$, we know $x^{1/6} \geq \ln(x)$ for any $x \geq 10^9$. $\square$

The next lemma shows quantitative properties of different variables under the setting of Theorem 4.

Lemma D.3. *Fix $\delta > 0$. Suppose that $T \geq T_1$ and $\beta = \left(\frac{\delta}{12(M+G)T}\right)^{1/4}$. We have $T \geq 36\Delta^{\mathrm{WCB}}$ and $\frac{378MG\ln^3(1/\beta)}{\varepsilon^2 T} \leq 6T(M+G)\beta^4$.*

Proof. To prove $T \geq 36\Delta^{\mathrm{WCB}}$, it suffices to show $\varepsilon^2 T \geq \frac{36\cdot 43}{\hat{s}_{\min}}\varepsilon G\sqrt{T}\ln(1/\beta)$ as $\sqrt{T} \geq \frac{1}{\varepsilon}$. Since $\beta = \left(\frac{\delta}{12(M+G)T}\right)^{1/4}$, we have $\ln(1/\beta) \leq \frac{1.25}{4}\ln(T/\delta) \leq \frac{5}{16}\ln(T)$ since $T \geq \delta^{-1}$. Therefore, it suffices to show $\varepsilon\sqrt{T} \geq \frac{1000G\ln T}{\hat{s}_{\min}}$. Note that by assumption, $T \geq 10^9$; as a result of Lemma D.2, $\ln T \leq T^{1/6}$. It remains to see $T^{1/3} \geq \frac{10^3 G}{\hat{s}_{\min}}$, which holds true because $T \geq T_1$.

We next show the second part of the result. Note that it is sufficient to show $T^2 \geq \frac{32(M+G)\ln^3(1/\beta)}{\varepsilon^2\beta^4}$ since $MG \leq \frac{1}{2}(M+G)^2$. Since $\beta = \left(\frac{\delta}{12T(M+G)}\right)^{1/4}$, we have $\frac{1}{\beta} = \left(\frac{12T(M+G)}{\delta}\right)^{1/4} \geq (12\cdot 10^9)^{1/4} \geq 100$. By Lemma D.2, we have $\ln^3(1/\beta) \leq \frac{1}{\beta}$ and thus to prove the desired result, it is sufficient to show $T^2 \geq \frac{32(M+G)}{\varepsilon^2\beta^5}$, which is implied by $T \geq \frac{6\sqrt{M+G}}{\varepsilon\beta^{2.5}}$. Plugging in the choice of β, it suffices to guarantee $T \geq \frac{6\sqrt{M+G}}{\varepsilon\left(\frac{\delta}{12T(M+G)}\right)^{2.5/4}}$ and it is implied by $T^{1/3} \geq \frac{72(M+G)^{4.5/4}}{\varepsilon\delta^{2.5/4}}$. By assumption, $T \geq \frac{10^9(M+G)^4}{\hat{s}_{\min}^3\varepsilon^3\delta^2}$, which satisfies the requirement. We thus complete the proof. $\square$

We are ready to prove Theorem 4.

Proof of Theorem 4. By Lemma D.3, $T \geq 36\Delta^{\mathrm{WCB}}$ and thus Lemma D.1 holds under the setting of β in the assumption of Theorem 4. As a result, with probability at least $C_1 \triangleq 1-(5T+5)(M+G)\beta^4 - \frac{378MG\ln^3(1/\beta)}{\varepsilon^2 T}$, we have $\mathcal{R}^{\mathrm{CBP}}_{g,\mathcal{F}} \leq C_2 \triangleq \frac{12\ln(1/\beta)}{\hat{s}_{\min}\varepsilon\sqrt{T}}\left(11\sqrt{G} + \frac{32G}{\varepsilon\sqrt{T}}\right)$. We next simplify C_1 and C_2 under the setting of Theorem 4. First for C_1, we know that $C_1 \geq 1 - 6T(M+G)\beta^4 - \frac{378MG\ln^3(1/\beta)}{\varepsilon^2 T}$ since $T > 5$. Lemma D.3 then implies that $C_1 \geq 1 - 12T(M+G)\beta^4$. Since $\beta = \left(\frac{\delta}{12(M+G)T}\right)^{1/4}$, we have $C_1 \geq 1-\delta$. Now for C_2, using the second inequality of (1) we immediately have $C_2 \leq \frac{200\ln(T/\delta)}{\hat{s}_{\min}\varepsilon}\sqrt{\frac{G}{T}}$. We thus finish the proof by plugging in the obtained lower bound of C_1 and the obtained upper bound of C_2. $\square$

D.2 Lower bound on T for high-probability g-regret (Assumption of Theorem 4)

Recall that in Theorem 4, we require $T \geq T_1$ where T_1 needs to scale with $1/\delta$ (see (10)). We next show that it is necessary by providing a lower bound of T to obtain high probability low g-regret result.

Proposition 3. *For any $\delta > 0$ and any even $T \leq \frac{0.15}{\delta}$, there exists an instance such that for any non-anticipatory algorithm, with probability at least δ, the g-regret of a fixed group is at least 0.15.*

Proof. Fix δ and an even $T \leq \frac{3}{40\delta}$. Consider a setting with $M = 2, G = 2, s_1 = \frac{T}{2}, p_1 = \frac{1}{T}, p_2 = 1 - \frac{1}{T}$ and we use the proportionally optimized fairness rule. For a case t from group 1, its score is given as follows: with probability $\frac{1}{2}$, $w_{t,1} = 1, w_{t,0} = 0$ and otherwise $w_{t,1} = 0, w_{t,2} = 1$. Then with probability $T(1 - {}^1\!/\!{}_T)^{T-1}\frac{1}{T} \geq 0.3$, there is exactly one case of group 1 and we have $O_1 = 0.5$. It also implies $\mathbb{E}[O_g] \geq 0.3 \times 0.5 = 0.15$. Now define event $\mathcal{S}$ by the event that there is only one case of group 1; this case is the last arrival; and the case has zero score for the only location J_1 with remaining capacity. Then $\mathbb{P}\{\mathcal{S}\} = \sum_{j=1}^{2} \mathbb{P}\{N(1, T-1) = 0, J_1 = j\}\mathbb{P}\{g(T) = 1, w_{T,j} = 0\} = \frac{0.5}{T}\mathbb{P}\{N(1, T-1) = 0\} \geq \frac{0.15}{T}$. In this scenario, $\alpha_g = 0$ for any non-anticipatory algorithm. Therefore, with probability at least $\frac{0.15}{T} \geq \delta$, it must incur $\alpha_g = 0 \leq \min(\mathbb{E}[O_g], O_g) - 0.15$. $\square$

D.3 Lower bound T_{emp} under Conservative Bid Price Control (Lemma 4.1)

Recall the definition of T_{emp} as the index of the first case after whose assignment a location runs out of free capacity. That is, $T_{\text{emp}} = \min\{t \colon \exists j \in \mathcal{M}, f_j(t+1) = 0\}$. Note that for the first T_{emp} cases, each arrival only uses free capacity (see Line 10 in Algorithm 2). Let $a_j(t)$ be the amount of consumed capacity at location j after the assignment of case t. We can equivalently write $T_{\text{emp}} = \min\{t \colon \exists j \in \mathcal{M}, a_j(t) = f_j(1) = s_j - \text{Res}(\beta)\}$. The intuition of the proof is to separate $a_j(t)$ into two parts: capacity used by assignment J^{WAB} and that used by greedy assignment. The first part is bounded using ideas similar to the proof of Lemma 3.3 for Amplified Bid Price Control. The second part is bounded by Lemma D.7.

To formalize the argument, let us define $\{b_j(t), t \in [T], j \in \mathcal{M}\}$ which concerns assignment J^{WAB}. For an arrival t, let $b_j(t) = 1$ if $J^{\text{WAB}}(t) = j$; otherwise $b_j(t) = 0$. To avoid the impact of capacity constraints, we consider a new system which has the same arrival process as the original one but constraints can be violated at no cost. Formally, define $\hat{J}^{\text{WCB}}(t)$ as the assignment for case t in the new system. Note that for the first T_{emp} cases, we have $J^{\text{CBP}}(t) = J^{\text{WCB}}(t) = \hat{J}^{\text{WCB}}(t)$ since the original system still has free capacity in every location. Based on $\hat{J}^{\text{WCB}}(t)$, we also similarly define $\hat{V}_g[t]$ by the total fairness excess score in the new system, given by $\sum_{\tau \in \mathcal{A}(g,t)}(w_{\tau, \hat{J}^{\text{WCB}}(\tau)} - \mathbb{E}[O_g])$. Note that condition (predict-to-meet) for $\hat{J}^{\text{WCB}}(t)$ is evaluated on $\hat{V}_g[t-1]$. Now let us define $c_g(t) \in \{0, 1\}$ such that $c_g(t) = 1$ if and only if case t is of group g and in the new system Algorithm 2 takes a greedy step, i.e., condition (predict-to-meet) is violated. Define $\hat{a}_j(t) = \sum_{\tau=1}^{t}\big(b_j(\tau) + c_{g(\tau)}(\tau)\big)$, which is an upper bound on the capacity usage in the new system. Also define $\widehat{T_{\text{emp}}} = \min\{t \colon \exists j \in \mathcal{M}, \hat{a}_j(t) \geq f_j(1)\}$. The next lemma shows that $\widehat{T_{\text{emp}}}$ is a lower bound of T_{emp}.

Lemma D.4. *For every sample path, we have $T_{\text{emp}} \geq \widehat{T_{\text{emp}}}$.*

Proof. Fix a case t' such that $a_j(t'-1) < f_j(1)$ for every $j \in \mathcal{M}$. That is, before the assignment of case t', every location has free capacity. As a result, for arrivals $1, \ldots, t'$, Conservative Bid

Price Control assigns cases following $J^{\mathrm{WCB}}(t')$, and thus $J^{\mathrm{CBP}}(t') = J^{\mathrm{WCB}}(t') = \hat{J}^{\mathrm{WCB}}(t')$. We then upper bound $a_j(t')$ for a location $j \in \mathcal{M}$ by

$$a_j(t') = \sum_{\tau=1}^{t'} \mathbb{1}\left(J^{\mathrm{CBP}}(\tau) = j\right) = \sum_{\tau=1}^{t'} \mathbb{1}\left(\hat{J}^{\mathrm{WCB}}(\tau) = j\right) \le \sum_{\tau=1}^{t'} (b_j(\tau) + c_{g(\tau)}(\tau)) = \hat{a}_j(t'), \tag{11}$$

where the inequality is due to the fact that as long as $\hat{J}^{\mathrm{WCB}}(\tau) = j$, either it is chosen because of an assignment J^{WAB} or a greedy assignment, i.e., $b_j(\tau) = 1$ or $c_{g(\tau)}(\tau) = 1$. Notice that for every $j \in \mathcal{M}$, we have $a_j(T_{\mathrm{emp}} - 1) < f_j(1)$ by the definition of T_{emp}. As a result, taking $t' = T_{\mathrm{emp}}$ in (11) gives $a_j(T_{\mathrm{emp}}) \le \hat{a}_j(T_{\mathrm{emp}})$ for every location j and thus $\hat{a}_{j'}(T_{\mathrm{emp}}) \ge f_{j'}(1)$ for some location j'. By the non-decreasing nature of $\hat{a}_{j'}(t)$ in t, we conclude that $T_{\mathrm{emp}} \ge \widehat{T_{\mathrm{emp}}}$. □

As a result of Lemma D.4, it suffices to show that with high probability $\widehat{T_{\mathrm{emp}}} \ge T - \Delta^{\mathrm{WCB}}$ to prove Lemma 4.1. Recall that $\widehat{T_{\mathrm{emp}}}$ is the first period t that $\hat{a}_j(t)$ exceeds $f_j(1)$ for some $j \in \mathcal{M}$. It thus requires us to upper bound $\hat{a}_j(t) = \sum_{\tau=1}^{t} b_j(\tau) + \sum_{\tau=1}^{t} c_{g(\tau)}(\tau)$. The first term on the right hand side corresponds to the number of assignments J^{WAB}. The second term is the number of greedy assignments. We upper bound these two terms separately in Lemma D.5 and Lemma D.7.

Lemma D.5. *With probability at least $1 - MT\beta^4$, we have $\sum_{\tau=1}^{t} b_j(\tau) \le \hat{s}_j t + \sqrt{2t\ln(1/\beta)}$ for every $t \in [T], j \in \mathcal{M}$.*

Proof. The proof is similar to that of Lemma 3.3. We include it here for completeness. Fix $t \in [T], j \in \mathcal{M}$. Note that $b_j(1), \ldots, b_j(t)$ are i.i.d. Bernoulli random variables. By Lemma C.1, we have $\mathbb{P}\{b_j(\tau) = 1\} = \mathbb{P}\{J^{\mathrm{WAB}}(\tau) = j\} \le \hat{s}_j = \frac{s_j}{T}$ for every $\tau \le t$. As a result, using Hoeffding's Inequality from Fact 3 gives $\mathbb{P}\left\{\sum_{\tau=1}^{t} b_j(\tau) \ge \hat{s}_j t + \sqrt{2t\ln(1/\beta)}\right\} \le \beta^4$. Using a union bound over all $t \in [T]$ and $j \in \mathcal{M}$ gives the desired result. □

Now let us upper bound the number of greedy steps. Instead of considering all groups, we count on a group level and show that for each group there are a limited number of greedy assignments during the entire horizon. Suppose the labels of group g cases are $t_1^g, \ldots, t_{N(g,T)}^g$. Condition on $N(g,T) = n$ where $1 \le n \le T$. Define $\mu_g^{\mathrm{GR}} = \mathbb{E}\left[w_{1,J^{\mathrm{GR}}(1)} \,\middle|\, g(1) = g\right], \mu_g^{\mathrm{WAB}} = \mathbb{E}\left[w_{1,J^{\mathrm{WAB}}(1)} \,\middle|\, g(1) = g\right]$ as the expected score of a group g arrival for a greedy and an assignment J^{WAB} respectively. In addition, let $n_g^{\mathrm{GR}}(t) = \sum_{\tau=1}^{t} c_g(\tau)$ be the number of greedy assignments for group g in the first t periods. The following lemma shows a lower bound on $\hat{V}_g[t_k^g]$ for any $1 \le k \le n$.

Lemma D.6. *For a fixed group g, there exists an event $\mathcal{S}_{\mathrm{score}}^g$ with probability at least $1 - 2T\beta^4$ such that condition on $N(g,T) = n$ with $n \in [1,T]$ and $\mathcal{S}_{\mathrm{score}}^g$, for any $1 \le k \le n$, we have*

$$\hat{V}_g[t_k^g] \ge n_g^{\mathrm{GR}}(t_k^g)\varepsilon_g - k\sqrt{\frac{1}{p_g T}} - 2\sqrt{k\ln(1/\beta)}.$$

Proof. Consider the sequence of scores for arrivals receiving greedy assignments W^{GR}. Define $\bar{\mu}_{g,i}^{\mathrm{GR}}$ as the average of the first i values in this sequence. Note that the assignment rule $\hat{J}^{\mathrm{WCB}}$ is independent of the score between case t and any location. As a result, elements in W^{GR} are i.i.d. and have the same distribution as $\max_j w_{1,j}$ condition on $g(1) = g$. Recall that $\mu_g^{\mathrm{GR}} = \mathbb{E}\left[\max_j w_{1,j} \,|\, g(1) = g\right]$. Then by Hoeffding's Inequality (Fact 3), for any fixed $i \le T$, we have $\mathbb{P}\left\{\bar{\mu}_{g,i}^{\mathrm{GR}} < \mu_g^{\mathrm{GR}} - \sqrt{\frac{2\ln(1/\beta)}{i}}\right\} \le \beta^4$. We can similarly define the sequence of scores for arrivals receiving assignment J^{WAB} by W^{WAB} and let $\bar{\mu}_{g,i}^{\mathrm{WAB}}$ be the average of the first i values. Also, recall that $\mu_g^{\mathrm{WAB}} = \mathbb{E}\left[w_{1,J^{\mathrm{WAB}}(1)} \,\middle|\, g(1) = g\right]$. Similar to the scenario of greedy scores, by Hoeffding's Inequality (Fact 3), for any fixed $1 \le i \le T$, we have $\mathbb{P}\left\{\bar{\mu}_{g,i}^{\mathrm{WAB}} < \mu_g^{\mathrm{WAB}} - \sqrt{\frac{2\ln(1/\beta)}{i}}\right\} \le \beta^4$. Define $\mathcal{S}_{\mathrm{score}}^g$ by

$$\mathcal{S}_{\mathrm{score}}^g = \left\{\forall i \le T,\ \bar{\mu}_{g,i}^{\mathrm{GR}} \ge \mu_g^{\mathrm{GR}} - \sqrt{\frac{2\ln(1/\beta)}{i}}\right\} \cap \left\{\forall i \le T,\ \bar{\mu}_{g,i}^{\mathrm{WAB}} \ge \mu_g^{\mathrm{WAB}} - \sqrt{\frac{2\ln(1/\beta)}{i}}\right\}$$

We know $\mathbb{P}\left\{\mathcal{S}_{\mathrm{score}}^g\right\} \ge 1 - 2T\beta^4$ by union bound. Let us now condition on $N(g,T) = n$. Fix k such that $1 \le k \le n$ and consider the kth arrival of group g. For ease of notation, let $n^{\mathrm{GR}} = n_g^{\mathrm{GR}}(t_k^g)$ and $n^{\mathrm{WAB}} = k - n^{\mathrm{GR}}$ be the number of greedy assignments and assignments J^{WAB} for group g cases in the first t_k^g periods. To prove the desired result, note that $n^{\mathrm{GR}} \le n$ and $n^{\mathrm{WAB}} \le n$. Therefore, under $\mathcal{S}_{\mathrm{score}}^g$, we have $\bar{\mu}_{n^{\mathrm{GR}}}^{\mathrm{GR}} \ge \mu_g^{\mathrm{GR}} - \sqrt{\frac{2\ln(1/\beta)}{n^{\mathrm{GR}}}}$ and $\bar{\mu}_{g,n^{\mathrm{WAB}}}^{\mathrm{WAB}} \ge \mu_g^{\mathrm{WAB}} - \sqrt{\frac{2\ln(1/\beta)}{n^{\mathrm{WAB}}}}$. By definition, $\hat{V}_g[t_c^g] = \sum_{i=1}^{n^{\mathrm{GR}}} s_i^{\mathrm{GR}} + \sum_{i=1}^{n^{\mathrm{WAB}}} s_i^{\mathrm{WAB}} - k\mathbb{E}\left[O_g\right] = \bar{\mu}_{g,n^{\mathrm{GR}}}^{\mathrm{GR}} n^{\mathrm{GR}} + \bar{\mu}_{g,n^{\mathrm{WAB}}}^{\mathrm{WAB}} n^{\mathrm{WAB}} - k\mathbb{E}\left[O_g\right]$. Then under $\mathcal{S}_{\mathrm{score}}$ we have

$$\begin{aligned}
\hat{V}_g[t_k^g] &= \bar{\mu}_{g,n^{\mathrm{GR}}}^{\mathrm{GR}} n^{\mathrm{GR}} + \bar{\mu}_{g,n^{\mathrm{WAB}}}^{\mathrm{WAB}} n^{\mathrm{WAB}} - k\mathbb{E}\left[O_g\right] \\
&\ge \mu_g^{\mathrm{GR}} n^{\mathrm{GR}} - \sqrt{2n^{\mathrm{GR}}\ln(1/\beta)} + \mu_g^{\mathrm{WAB}} n^{\mathrm{WAB}} - \sqrt{2n^{\mathrm{WAB}}\ln(1/\beta)} - k\mathbb{E}\left[O_g\right] \\
&\ge n^{\mathrm{GR}}\left(\mu_g^{\mathrm{GR}} - \mathbb{E}\left[O_g\right]\right) + (k - n^{\mathrm{GR}})\left(\mu_g^{\mathrm{WAB}} - \mathbb{E}\left[O_g\right]\right) - 2\sqrt{k\ln(1/\beta)} \\
&\ge n^{\mathrm{GR}}\varepsilon_g - k\sqrt{\frac{1}{p_g T}} - 2\sqrt{k\ln(1/\beta)} \quad = \quad n_g^{\mathrm{GR}}(t_k^g)\varepsilon_g - k\sqrt{\frac{1}{p_g T}} - 2\sqrt{k\ln(1/\beta)}
\end{aligned}$$

where the first inequality uses the definition of $\mathcal{S}_{\mathrm{score}}^g$; the second inequality uses $n^{\mathrm{GR}} + n^{\mathrm{WAB}} = k$ and $\sqrt{a} + \sqrt{b} \le \sqrt{2(a+b)}$ for any non-negative a, b; the last inequality uses the slackness property (Definition 2) that $\mu_g^{\mathrm{GR}} \ge \mathbb{E}\left[O_g\right] + \varepsilon_g$ and Lemma 3.2 that $\mu_g^{\mathrm{WAB}} \ge \mathbb{E}\left[O_g\right] - \sqrt{\frac{1}{p_g T}}$. $\square$

We can then upper bound the total number of greedy steps.

Lemma D.7. *We have* $\mathbb{P}\left\{\sum_{g \in \mathcal{G}} \sum_{\tau=1}^{T} c_g(\tau) \le \frac{9\sqrt{GT}\ln(1/\beta)}{\varepsilon} + \frac{25G\ln(1/\beta)}{\varepsilon^2}\right\} \ge 1 - (2T+1)G\beta^4$.

Proof of Lemma D.7. Fix a group g. We want to bound the number of greedy assignments. Condition on the number of group g arrivals by n, i.e., condition on $N(g,T) = n \in [T]$. Take $k \in [0, n-1]$ such that the $(k+1)$th arrival is the last case that receives a greedy assignment in group g. If there is no such arrival, we know $\sum_{\tau=1}^{T} c_g(t) = 0$. By definition of k, we have $\sum_{\tau=1}^{T} c_g(t) = n_g^{\mathrm{GR}}(t_k^g) + 1$. Therefore, it suffices to upper bound $n_g^{\mathrm{GR}}(t_k^g)$. Recall $\mathcal{S}_{\mathrm{score}}^g$ the event in Lemma D.6. We know $\mathbb{P}\{\mathcal{S}_{\mathrm{score}}^g\} \ge 1 - 2T\beta^4$. In addition, under $\mathcal{S}_{\mathrm{score}}^g$, we have

$$\hat{V}_g[t_k^g] \ge n_g^{\mathrm{GR}}(t_k^g)\varepsilon_g - k\sqrt{\frac{1}{p_g T}} - 2\sqrt{k\ln(1/\beta)} \ge n_g^{\mathrm{GR}}(t_k^g)\varepsilon_g - n\sqrt{\frac{1}{p_g T}} - 2\sqrt{n\ln(1/\beta)}. \tag{12}$$

Note that since arrival $k+1$ receives a greedy assignment, by the definition of J^{WCB} (Line 5 of Algorithm 2) and the setting of parameters in (PAR), we have $\hat{V}_g[t_k^g] \le \mathrm{Buf}(\beta) + \Psi(g,t,\beta) + 1 = 6\ln(1/\beta) + \frac{\ln(1/\beta)}{2\varepsilon_g} + 1 \le \frac{8\ln(1/\beta)}{\varepsilon_g}$. Combining this with (12) implies $n_g^{\mathrm{GR}}(t_k^g)\varepsilon_g - n\sqrt{\frac{1}{p_g T}} - 2\sqrt{n\ln(1/\beta)} \le \frac{8\ln(1/\beta)}{\varepsilon_g}$. As a result, condition on $N(g,T)=n$ and $\mathcal{S}^g_{\mathrm{score}}$, we have

$$\sum_{\tau=1}^{T} c_g(\tau) = n_g^{\mathrm{GR}}(t_k^g) + 1 \le \frac{n}{\sqrt{p_g T}\varepsilon_g} + \frac{2\sqrt{n\ln(1/\beta)}}{\varepsilon_g} + \frac{8\ln(1/\beta)}{\varepsilon_g^2} + 1 \le \frac{n}{\sqrt{p_g T}\varepsilon_g} + \frac{2\sqrt{n\ln(1/\beta)}}{\varepsilon_g} + \frac{9\ln(1/\beta)}{\varepsilon_g^2}. \tag{13}$$

Define $\bar{n} = p_g T + 5\left(\sqrt{p_g T} + \ln(1/\beta)\right)$ and $\mathcal{S}^g_{\mathrm{size}} = \{N(g,T) \le \bar{n}\}$. Using Lemma E.2, we know that $\mathbb{P}\{\mathcal{S}^g_{\mathrm{size}}\} \ge 1-\beta^4$. Condition on both $\mathcal{S}^g_{\mathrm{size}}$ and $\mathcal{S}^g_{\mathrm{score}}$. If $p_g T \le 1$, then $\sum_{t=1}^{T} c_g(t) \le \bar{n} \le 6 + 5\ln(1/\beta)$. Otherwise, using (13) gives

$$\begin{aligned}
\sum_{\tau=1}^{T} c_g(\tau) &\le \frac{\bar{n}}{\sqrt{p_g T}\varepsilon_g} + \frac{2\sqrt{\bar{n}\ln(1/\beta)}}{\varepsilon_g} + \frac{9\ln(1/\beta)}{\varepsilon_g^2} \\
&= \frac{p_g T + 5\left(\sqrt{p_g T\ln(1/\beta)} + \ln(1/\beta)\right)}{\sqrt{p_g T}\varepsilon_g} + \frac{2\sqrt{\left(p_g T + 5\left(\sqrt{p_g T\ln(1/\beta)} + \ln(1/\beta)\right)\right)\ln(1/\beta)}}{\varepsilon_g} \\
&\quad + 9\ln(1/\beta)/\varepsilon_g^2 \\
&\le \frac{\sqrt{p_g T}}{\varepsilon_g} + \frac{10\ln(1/\beta)}{\varepsilon_g} + \frac{2\sqrt{p_g T\ln(1/\beta)}}{\varepsilon_g} + \frac{6\left(p_g T\ln^3\frac{1}{\beta}\right)^{0.25}}{\varepsilon_g} + \frac{6\ln(1/\beta)}{\varepsilon_g} + \frac{9\ln(1/\beta)}{\varepsilon_g^2} \\
&\le \frac{9\sqrt{p_g T\ln(1/\beta)}}{\varepsilon_g} + \frac{25\ln(1/\beta)}{\varepsilon_g^2}
\end{aligned} \tag{14}$$

where the second to last inequality uses the assumption that $p_g T \ge 1$ and the fact that $\sqrt{a+b} \le \sqrt{a} + \sqrt{b}$ for any $a, b \ge 0$. Combining the two scenarios (whether $p_g T \le 1$ or not), for every $g \in \mathcal{G}$, we have $\sum_{\tau=1}^{T} c_g(\tau) \le \frac{9\sqrt{p_g T\ln(1/\beta)}}{\varepsilon_g} + \frac{25\ln(1/\beta)}{\varepsilon_g^2}$.

Finally, let us define $\mathcal{S} = \cap_{g\in\mathcal{G}}\left(\mathcal{S}^g_{\mathrm{size}} \cap \mathcal{S}^g_{\mathrm{score}}\right)$. Then by union bound, we have $\mathbb{P}\{\mathcal{S}\} \ge 1 - \sum_{g\in\mathcal{G}}(1 - \mathbb{P}\{\mathcal{S}^g_{\mathrm{size}}\} + 1 - \mathbb{P}\{\mathcal{S}^g_{\mathrm{score}}\}) \ge 1 - G(2T+1)\beta^4$. In addition, condition on $\mathcal{S}$, the total number of greedy assignments is upper bounded by

$$\begin{aligned}
\sum_{g\in\mathcal{G}}\sum_{\tau=1}^{T} c_g(\tau) \le \sum_{g\in\mathcal{G}}\left(\frac{9\sqrt{p_g T\ln(1/\beta)}}{\varepsilon_g} + \frac{25\ln(1/\beta)}{\varepsilon_g^2}\right) &\le \sum_{g\in\mathcal{G}}\left(\frac{9\sqrt{p_g T\ln(1/\beta)}}{\varepsilon} + \frac{25\ln(1/\beta)}{\varepsilon^2}\right) \\
&\le \frac{9\sqrt{GT\ln(1/\beta)}}{\varepsilon} + \frac{25G\ln(1/\beta)}{\varepsilon^2}.
\end{aligned}$$

The first inequality is shown in (14); the second inequality follows from $\varepsilon = \min_g \varepsilon_g$; the third inequality follows from the Cauchy-Schwarz inequality since $\sum_{g\in\mathcal{G}}\sqrt{p_g} \le \sqrt{G\left(\sum_{g\in\mathcal{G}} p_g\right)} = \sqrt{G}$. $\square$

Proof of Lemma 4.1. Recall the definition of $\widehat{T_{\mathrm{emp}}}$ as the first period t that $\hat{a}_j(t)$ is at least the initial free capacity. That is, $\widehat{T_{\mathrm{emp}}} = \min\{t\colon \exists j \in \mathcal{M}, \hat{a}_j(t) \ge f_j(1)\}$. By definition and (PAR), we have $f_j(1) = s_j - G\mathrm{Cap}(\beta) \ge s_j - 7G\ln(1/\beta)$. Using union bound over the events in Lemma D.5 and

Lemma D.7, there is an event $\mathcal{S}$ with $\mathbb{P}\{\mathcal{S}\} \geq 1 - MT\beta^4 - (2T+1)G\beta^4 \geq 1-(2T+1)(M+G)\beta^4$, such that under $\mathcal{S}$, we have $\sum_{g\in\mathcal{G}}\sum_{\tau=1}^T c_g(\tau) \leq \frac{9\sqrt{GT}\ln(1/\beta)}{\varepsilon} + \frac{25G\ln(1/\beta)}{\varepsilon^2}$ and that for every $g \in \mathcal{G}, t\in[T]$, we have $\sum_{\tau=1}^t b_j(\tau) \leq \hat{s}_j t + \sqrt{2t\ln(1/\beta)}$. As a result, under $\mathcal{S}$, for every $t\in[T], j\in\mathcal{M}$,

$$\begin{aligned}\hat{a}_j(t) \leq \sum_{\tau=1}^t \big(b_j(\tau) + c_{g(\tau)}(\tau)\big) &\leq \sum_{\tau=1}^t b_j(\tau) + \sum_{g\in\mathcal{G}}\sum_{\tau=1}^T c_g(\tau)\\ &\leq \hat{s}_j t + \sqrt{2t\ln(1/\beta)} + \frac{9\sqrt{GT}\ln(1/\beta)}{\varepsilon} + \frac{25G\ln(1/\beta)}{\varepsilon^2}.\end{aligned}$$

Recall $\Delta^{\mathrm{WCB}} = (\min_{j\in\mathcal{M}} \hat{s}_j)^{-1}\left(\frac{11\sqrt{GT}\ln(1/\beta)}{\varepsilon} + \frac{32G\ln(1/\beta)}{\varepsilon^2}\right) > 0$. We have $\forall j\in\mathcal{M}$,

$$\begin{aligned}\hat{a}_j(T-\Delta^{\mathrm{WCB}}) &\leq \hat{s}_j(T-\Delta^{\mathrm{WCB}}) + \sqrt{2(T-\Delta^{\mathrm{WCB}})\ln(1/\beta)} + \frac{9\sqrt{GT}\ln(1/\beta)}{\varepsilon} + \frac{25G\ln(1/\beta)}{\varepsilon^2}\\ &< \hat{s}_j(T-\Delta^{\mathrm{WCB}}) + \sqrt{2T\ln(1/\beta)} + \frac{9\sqrt{GT}\ln(1/\beta)}{\varepsilon} + \frac{25G\ln(1/\beta)}{\varepsilon^2}\\ &\leq \hat{s}_j T - \hat{s}_j\Delta^{\mathrm{WCB}} + \frac{11\sqrt{GT}\ln(1/\beta)}{\varepsilon} + \frac{25G\ln(1/\beta)}{\varepsilon^2} \leq s_j - \frac{7G\ln(1/\beta)}{\varepsilon^2}\end{aligned}$$

where the third inequality uses $\sqrt{2T\ln(1/\beta)} \leq \frac{2\sqrt{T}\ln(1/\beta)}{\varepsilon}$ and the last inequality uses the definition of Δ^{WCB}. Therefore, for every location j, we have $\hat{a}_j(T-\Delta^{\mathrm{WCB}}) < s_j - \frac{7G\ln(1/\beta)}{\varepsilon^2} \leq f_j(1)$. This implies that $\widehat{T_{\mathrm{emp}}} \geq T - \Delta^{\mathrm{WCB}}$ under $\mathcal{S}$. We then complete the proof by using Lemma D.4 that $T_{\mathrm{emp}} \geq \widehat{T_{\mathrm{emp}}}$ and that $\mathbb{P}\{\mathcal{S}\} \geq 1-(2T+1)(M+G)\beta^4$. □

D.4 Predict-to-meet condition ensures low g-regret (Lemma 4.2)

We first show that under the setting of parameters in (PAR), $\Psi(g,t,\beta)$ is a valid lower bound (with high probability) on future fairness score surplus for group g assuming all remaining arrivals after case t obtain greedy selections.

Lemma D.8. *There exists an event $\mathcal{S}_{\mathrm{low}}$ with probability at least $1-3GT\beta^4$ such that: $\forall g\in\mathcal{G}, t\in[T]$, we have $\Psi(g,t,\beta) \leq \sum_{\tau\in\mathcal{A}(g,T)\setminus\mathcal{A}(g,t)}\left(w_{\tau,J^{\mathrm{GR}}(\tau)} - \mathbb{E}\left[O_g\right]\right)$.*

Proof. Fix a group g. Since the number of cases T is known, greedy scores of arrivals from group g counting backward from the end of the horizon can be viewed as a sequence of i.i.d. random variables. By Hoeffding's Inequality (Fact 3) and union bound over $1,\ldots,T$, we have with probability at least $1-T\beta^4$ that for every $0\leq n\leq T$, the sum of greedy scores for the last n arrivals from group g is lower bounded by $n\mathbb{E}\left[w_{1,J^{\mathrm{GR}}(1)}\,\middle|\, g(1)=g\right] - \sqrt{2n\ln(1/\beta)}$. Denote this event by $\mathcal{S}_{\mathrm{low}}^g$. Note that for every $t\in[T]$, the number of group g arrivals between $t+1$ and T is at most T. For ease of notation, let us write $N^f(g,t) = N(g,T) - N(g,t)$ for the number of "future" cases of group g after case t. Then under $\mathcal{S}_{\mathrm{low}}^g$, for every $t\in[T]$, we have

$$\sum_{\tau=t+1}^T \mathbb{1}\left(g(\tau)=g\right) w_{\tau,J^{\mathrm{GR}}(\tau)} \geq N^f(g,t)\mathbb{E}\left[w_{1,J^{\mathrm{GR}}(1)}\,\middle|\, g(1)=g\right] - \sqrt{2N^f(g,t)\ln(1/\beta)}. \tag{15}$$

Note that by definition $\mathbb{E}\left[w_{1,J^{\mathrm{GR}}(1)} \,\middle|\, g(1)=g\right] = \mathbb{E}\left[O_g\right] + \varepsilon_g$, under $\mathcal{S}^g_{\mathrm{low}}$, (15) implies

$$\sum_{\tau=t+1}^{T} \mathbb{1}\left(g(\tau)=g\right)\left(w_{\tau,J^{\mathrm{GR}}(\tau)} - \mathbb{E}\left[O_g\right]\right) \geq N^f(g,t)\varepsilon_g - \sqrt{2N^f(g,t)\ln(1/\beta)}. \tag{16}$$

We next lower bound the right hand side. Let $f(n) = n\varepsilon_g - \sqrt{2n\ln(1/\beta)}$. We know $f(0)=0$. Taking derivative of $f(n)$ gives $f'(n) = \varepsilon_g - \sqrt{\frac{\ln(1/\beta)}{2n}}$. Setting the derivative to zero gives the unique stationary point $n^\star = \frac{\ln(1/\beta)}{2\varepsilon_g^2}$. We then know that $\forall n \geq 0, f(n) \geq -\frac{\ln(1/\beta)}{2\varepsilon_g}$. Since $N^f(g,t) \geq 0$, we have $\sum_{\tau=t+1}^{T} \mathbb{1}\left(g(\tau)=g\right)\left(w_{\tau,J^{\mathrm{GR}}(\tau)} - \mathbb{E}\left[O_g\right]\right) \geq -\frac{\ln(1/\beta)}{2\varepsilon_g} = \Psi(g,t,\beta)$. We then finish the proof by using the union bound over all groups and observing that the event $\mathcal{S}_{\mathrm{low}}$ defined by $\{\cap_{g\in\mathcal{G}}\mathcal{S}^g_{\mathrm{low}}\}$ occurs with probability at least $1-GT\beta^4$. $\square$

Proof of Lemma 4.2. For a group g, let us define an event $\mathcal{S}_g$ under which the number of group g arrivals in $[1,T]$ is close to its expectation (from below) and that in $[1, T-\Delta^{\mathrm{WCB}}]$ is close to its expectation from above. In particular, $\mathcal{S}_g$ is the intersection of the event that $N(g,T) \geq p_gT - 3\sqrt{p_gT\ln(\frac{1}{\beta})}$ and the event that $N(g,T) - N(g,T-\Delta^{\mathrm{WCB}}) \leq p_g\Delta^{\mathrm{WCB}} + 5\sqrt{\max\left(p_g\Delta^{\mathrm{WCB}}, \ln(\frac{1}{\beta})\right)\ln(\frac{1}{\beta})}$. Then by concentration bound (Lemma E.2), we have $\mathbb{P}\{\mathcal{S}_g\} \geq 1-2\beta^4$. Let us define $\mathcal{S}_{\mathrm{BP}}$ by

$$\mathcal{S}_{\mathrm{BP}} = \{T_{\mathrm{emp}} \geq T - \Delta^{\mathrm{WCB}}\} \cap \mathcal{S}_{\mathrm{low}} \cap \left(\cap_{g\in\mathcal{G}}\mathcal{S}_g\right) \tag{17}$$

where $\mathcal{S}_{\mathrm{low}}$ is defined in Lemma D.8. By Lemma 4.1 and Lemma D.8, we have $\mathbb{P}\{T_{\mathrm{emp}} \geq T - \Delta^{\mathrm{WCB}}\} \geq 1-(2T+1)(M+G)\beta^4$ and $\mathbb{P}\{\mathcal{S}_{\mathrm{low}}\} \geq 1-GT\beta^4$. Using union bound gives $\mathbb{P}\{\mathcal{S}_{\mathrm{BP}}\} \geq 1-(3T+1)(M+G)\beta^4 - 2G\beta^4 \geq 1-(3T+3)(M+G)\beta^4$. We next show that condition on $\mathcal{S}_{\mathrm{BP}}$, we have the desired properties in Lemma 4.2. We first have $T_{\mathrm{emp}} \geq T-\Delta^{\mathrm{WCB}}$ since $\mathcal{S}_{\mathrm{BP}} \subseteq \{T_{\mathrm{emp}} \geq T-\Delta^{\mathrm{WCB}}\}$.

Let us then fix a group $g \in \mathcal{G}$. Suppose that condition predict-to-meet is satisfied for case t and the case is of group g. If there are multiple such cases, we take t to be the label of the last one. Our goal is to show that condition on $\mathcal{S}_{\mathrm{BP}}$ (17) and assuming the existence of t, we have $\alpha_g \geq \mathbb{E}\left[O_g\right] - \frac{12\Delta^{\mathrm{WCB}}}{T}$. By condition predict-to-meet, we have $V_g[t-1] - \mathbb{E}\left[O_g\right] + \Psi(g,t,\beta) \geq \mathrm{Buf}(\beta)$. Then by the definition of $V_g[t-1]$, the fact that $\mathcal{S}_{\mathrm{BP}} \subseteq \mathcal{S}_{\mathrm{low}}$, Lemma D.8 and (PAR), we have

$$\sum_{\tau\in\mathcal{A}(g,t-1)}\left(w_{\tau,J^{\mathrm{CBP}}(\tau)} - \mathbb{E}\left[O_g\right]\right) - \mathbb{E}\left[O_g\right] + \sum_{\tau=t+1}^{T}\mathbb{1}\left(g(\tau)=g\right)\left(w_{\tau,J^{\mathrm{GR}}}(\tau) - \mathbb{E}\left[O_g\right]\right) \geq 6\ln(1/\beta), \tag{18}$$

which by reorganizing terms implies

$$\sum_{\tau\in\mathcal{A}(g,t)} w_{\tau,J^{\mathrm{CBP}}(\tau)} + \sum_{\tau=t+1}^{T}\mathbb{1}\left(g(\tau)=g\right) w_{\tau,J^{\mathrm{GR}}}(\tau) - 6\ln(1/\beta) \geq N(g,T)\mathbb{E}\left[O_g\right].$$

Recall that $\alpha_g = \frac{1}{N(g,T)}\sum_{\tau\in\mathcal{A}(g,T)} w_{\tau,J^{\mathrm{CBP}}(\tau)}$. Using the definition of α_g and (18) gives

$$\alpha_g \geq \mathbb{E}\left[O_g\right] - \frac{1}{N(g,T)}\left(\sum_{\tau=t+1}^{T}\mathbb{1}\left(g(\tau)=g\right)\left(w_{\tau,J^{\mathrm{GR}}(\tau)} - w_{\tau,J^{\mathrm{CBP}}(\tau)}\right) - 6\ln(1/\beta)\right). \tag{19}$$

Note that we have $N(g,T) \geq 1$ due to the existence of t. It remains to show that the second term on the right hand side can be upper bounded by $\frac{12\Delta^{\mathrm{WCB}}}{T}$. Notice that $T_{\mathrm{emp}} \geq T - \Delta^{\mathrm{WCB}}$ and t is the last case of group g for whom condition (predict-to-meet) is satisfied. We next consider two scenarios. On the one hand, if $t \geq T - \Delta^{\mathrm{WCB}}$, we must have $\sum_{\tau=t+1}^{T} \mathbb{1}\left(g(\tau) = g\right)\left(w_{\tau,J^{\mathrm{GR}}(\tau)} - w_{\tau,J^{\mathrm{CBP}}(\tau)}\right) \leq N(g,T) - N(g, T-\Delta^{\mathrm{WCB}})$ since scores are in $[0,1]$. On the other hand, suppose that $t < T - \Delta^{\mathrm{WCB}}$. Under $\mathcal{S}_{\mathrm{BP}}$, we have $T_{\mathrm{emp}} \geq T - \Delta^{\mathrm{WCB}}$ and thus $J^{\mathrm{CBP}}(\tau) = J^{\mathrm{WCB}}(\tau)$ for every $\tau \in [t, T - \Delta^{\mathrm{WCB}}]$. In addition, since t is the last case of group g such that condition (predict-to-meet) is satisfied, we have $J^{\mathrm{CBP}}(\tau) = J^{\mathrm{WCB}}(\tau) = J^{\mathrm{GR}}(\tau)$ for every $\tau \in [t+1, T-\Delta^{\mathrm{WCB}}]$ with $g(\tau) = g$. As a result, under the scenario that $t < T - \Delta^{\mathrm{WCB}}$, we again have

$$\begin{aligned}\sum_{\tau=t+1}^{T} \mathbb{1}\left(g(\tau) = g\right)\left(w_{\tau,J^{\mathrm{GR}}(\tau)} - w_{\tau,J^{\mathrm{CBP}}(\tau)}\right) &= \sum_{\tau=T-\Delta^{\mathrm{WCB}}+1}^{T} \mathbb{1}\left(g(\tau) = g\right)\left(w_{\tau,J^{\mathrm{GR}}(\tau)} - w_{\tau,J^{\mathrm{CBP}}(\tau)}\right)\\ &\leq N(g,T) - N(g,T-\Delta^{\mathrm{WCB}}).\end{aligned}$$

Summarizing the two scenarios and with (19), we have $\alpha_g \geq \mathbb{E}\left[O_g\right] - \frac{N(g,T)-N(g,T-\Delta^{\mathrm{WCB}})-6\ln(1/\beta)}{N(g,T)}$. Recall the definition of $\mathcal{S}_g$ and the fact that $\mathcal{S}_{\mathrm{BP}} \subseteq \mathcal{S}_g$ by (17). We have that $N(g,T) \geq p_g T - 3\sqrt{p_g T \ln(1/\beta)}$ and that $N(g,T) - N(g,T-\Delta^{\mathrm{WCB}}) \leq p_g\Delta^{\mathrm{WCB}} + 5\sqrt{\max\left(p_g\Delta^{\mathrm{WCB}}, \ln(1/\beta)\right)\ln(1/\beta)}$. Consider two scenarios on the size of group g. First, if $p_g\Delta^{\mathrm{WCB}} \leq \ln(1/\beta)$, we have that $N(g,T) - N(g,T-\Delta^{\mathrm{WCB}}) \leq \ln(1/\beta) + 5\ln(1/\beta) \leq 6\ln(1/\beta)$. Therefore, for a group g with $p_g\Delta^{\mathrm{WCB}} \leq \ln(1/\beta)$,

$$\alpha_g \geq \mathbb{E}\left[O_g\right] - \frac{N(g,T) - N(g,T-\Delta^{\mathrm{WCB}}) - 6\ln(1/\beta)}{N(g,T)} \geq \mathbb{E}\left[O_g\right] - \frac{6\ln(1/\beta) - 6\ln(1/\beta)}{N(g,T)} = \mathbb{E}\left[O_g\right].$$

Otherwise, suppose that $p_g\Delta^{\mathrm{WCB}} \geq \ln(1/\beta)$. We then have $N(g,T) - N(g,T-\Delta^{\mathrm{WCB}}) \leq 6p_g\Delta^{\mathrm{WCB}}$. In addition,

$$N(g,T) \geq p_gT - 3\sqrt{p_gT\ln(1/\beta)} \geq p_gT - 3\sqrt{p_gT\left(p_g\Delta^{\mathrm{WCB}}\right)} \geq p_gT - 3\sqrt{p_gT\left(\frac{1}{36}p_gT\right)} = \frac{1}{2}p_gT,$$

where the last inequality is by assumption that $T \geq 36\Delta^{\mathrm{WCB}}$. As a result, when $p_g\Delta^{\mathrm{WCB}} \geq \ln(1/\beta)$, we have $\alpha_g \geq \mathbb{E}\left[O_g\right] - \frac{N(g,T)-N(g,T-\Delta^{\mathrm{WCB}})}{N(g,T)} \geq \frac{6p_g\Delta^{\mathrm{WCB}}}{0.5p_gT} = \frac{12\Delta^{\mathrm{WCB}}}{T}$. Summarizing the above two scenarios shows that condition on $\mathcal{S}_{\mathrm{BP}}$, for every group g, if there is an arrival satisfying condition (predict-to-meet), we have that $\alpha_g \geq \min\left(\mathbb{E}\left[O_g\right], \mathbb{E}\left[O_g\right] - \frac{12\Delta^{\mathrm{WCB}}}{T}\right) = \mathbb{E}\left[O_g\right] - \frac{12\Delta^{\mathrm{WCB}}}{T}$. □

D.5 Guarantee of g-regret for large groups (Lemma 4.3)

Proof. Fix a group g with $p_g \geq \frac{54\ln^2(1/\beta)}{\varepsilon_g^2 T}$. Our proof involves the following steps: 1) we show that there are sufficiently many arrivals of group g before $T - \Delta^{\mathrm{WCB}}$; 2) we show that under this event, the total greedy score is so high that Conservative Bid Price Control must assign at least one arrival of group g before case $T - \Delta^{\mathrm{WCB}}$ to J^{WAB}; 3) we apply Lemma 4.2 for the final result.

We first show there are sufficient arrivals of group g before $T - \Delta^{\mathrm{WCB}}$. Let $n^\star = \frac{12\ln^2(1/\beta)}{\varepsilon_g^2}$. Define $\mathcal{S}_1 = \{N(g, T-\Delta^{\mathrm{WCB}}) \geq n^\star + 1\}$. Note that by assumption, $T \geq 36\Delta^{\mathrm{WCB}}$ and thus we know that $\mathbb{E}\left[N(g,T-\Delta^{\mathrm{WCB}})\right] = p_g(T - \Delta^{\mathrm{WCB}}) \geq \frac{35}{36}p_gT \geq \ln(1/\beta)$. Using the last probability bound of Lemma E.2, with probability at least $1-\beta^4$, we have $N(g,T-\Delta^{\mathrm{WCB}}) \geq p_g(T-\Delta^{\mathrm{WCB}}) -$

$5\sqrt{p_gT\ln(1/\beta)}$. Since $p_gT \geq 54\ln^2(1/\beta)$, we have $5\sqrt{p_gT\ln(1/\beta)} \leq \frac{5}{\sqrt{54}}p_gT \leq \frac{5}{7}p_gT$. Therefore, with probability at least $1-\beta^4$, $N(g, T-\Delta^{\text{WCB}})$ is lower bounded by

$$p_g(T-\Delta^{\text{WCB}})-5\sqrt{p_gT\ln(1/\beta)} \geq \frac{35}{36}p_gT-\frac{5}{7}p_gT \geq \frac{35}{36}p_gT-\frac{3}{4}p_gT+1 = \frac{2}{9}p_gT+1 \geq \frac{12\ln^2(1/\beta)}{\varepsilon_g^2}+1.$$

Therefore, $\mathbb{P}\{\mathcal{S}_1\} \geq 1-\beta^4$.

We next define $\mathcal{S}_2$ as the event that the total greedy score of the first $n^\star$ arrivals of group g is at least $n^\star\mathbb{E}\left[w_{1,J^{\text{GR}}(1)}\,\middle|\,g(1)=g\right]-\sqrt{2n^\star\ln(1/\beta)}$. Since scores of arrivals are i.i.d., using Hoeffding's inequality (Fact 3) gives $\mathbb{P}\{\mathcal{S}_2\} \geq 1-\beta^4$. Recall event $\mathcal{S}_{\text{BP}}$ in Lemma 4.2. Under $\mathcal{S}_{\text{BP}}$, we have $T_{\text{emp}} \geq T-\Delta^{\text{WCB}}$ and that if a group g receives an assignment J^{WAB} during $[1, T_{\text{emp}}]$, we have $\alpha_g \geq \mathbb{E}[O_g]-\frac{12\Delta^{\text{WCB}}}{T}$. Condition on $\mathcal{S}_1\cap\mathcal{S}_2\cap\mathcal{S}_{\text{BP}}$. It remains to show that there exists a case $t \leq T_{\text{emp}}$, such that $g(t)=g$ and condition (predict-to-meet) holds (so group g receives an assignment J^{WAB}.) Note that conditioned on $\mathcal{S}_{\text{BP}}$, we have $T_{\text{emp}} \geq T-\Delta^{\text{WCB}}$. In addition, by $\mathcal{S}_1$, the number of group g arrivals in the first $T-\Delta^{\text{WCB}}$ periods is at least $n^\star+1$. Consider two scenarios. First, assume that for the first $n^\star$ group g arrivals, condition (predict-to-meet) holds at least once. Then we know there must exist $t' \leq T-\Delta^{\text{WCB}} \leq T_{\text{emp}}$ with $g(t')=g$ satisfying condition (predict-to-meet). Otherwise, the first $n^\star$ group g cases all receive their greedy assignments. Let t' be the label of the $n^\star+1$th cases of group g. We next show condition (predict-to-meet) holds for case t'. We first have $t' \leq T-\Delta^{\text{WCB}}$ since $N(g, T-\Delta^{\text{WCB}}) \geq n^\star+1$. The definition of $V_g[t'-1]$ gives $V_g[t'-1] = \sum_{\tau<t',\ g(\tau)=g}(w_{\tau,J^{\text{CBP}}(\tau)}-\mathbb{E}[O_g])$. We then have

$$\begin{aligned} V_g[t'-1] &= \sum_{\tau\in\mathcal{A}(g,t'-1)} w_{\tau,J^{\text{GR}}(\tau)}-n^\star\mathbb{E}[O_g] \\ &\geq n^\star\mathbb{E}\left[w_{1,J^{\text{GR}}(1)}\,\middle|\,g(1)=g\right]-\sqrt{2n^\star\ln(1/\beta)}-n^\star\mathbb{E}[O_g] \geq n^\star\epsilon_g-\sqrt{2n^\star\ln(1/\beta)}, \end{aligned}$$

where the first equality is because the first $n^\star$ cases of group g receive greedy assignments; the second inequality is by event $\mathcal{S}_2$; the third inequality is by the slackness property (Definition 2). By the definition of $n^\star = \frac{12\ln^2(1/\beta)}{\varepsilon_g^2}$, we have

$$\begin{aligned} V_g[t'-1]-\left(6\ln(1/\beta)+\frac{\ln(1/\beta)}{2\varepsilon_g}+1\right) &\geq \frac{12\ln^2(1/\beta)}{\varepsilon_g}-6\ln(1/\beta)-\frac{2\sqrt{6}\ln(1/\beta)}{\varepsilon_g}-\frac{\ln(1/\beta)}{\varepsilon_g} \\ &\geq \frac{\ln(1/\beta)}{\varepsilon}\left(12-6-2\sqrt{6}-1\right) \geq 0. \end{aligned}$$

Under the parameter setting in (PAR), for case t', condition (predict-to-meet) holds. Therefore, condition on $\mathcal{S}_1\cap\mathcal{S}_2\cap\mathcal{S}_{\text{BP}}$, we have that group g receives an assignment J^{WAB} during $[1, T_{\text{emp}}]$ and $\alpha_g \geq \mathbb{E}[O_g]-\frac{12\Delta^{\text{WCB}}}{T}$ by Lemma 4.2. We then finish the proof using $\mathbb{P}\{\mathcal{S}_{\text{BP}}\} \geq 1-(5T+1)(M+G)\beta^4$ and noticing that

$$\mathbb{P}\{S_1\cap S_2\cap\mathcal{S}_{\text{BP}}\} \geq 1-(5T+1)(M+G+1)\beta^4-2\beta^4 \geq 1-(5T+3)(M+G+1)\beta^4.$$

□

D.6 Guarantee of g-regret for small groups (Lemma 4.4)

Proof. Recall event $\mathcal{S}_{\text{BP}}$ defined in Lemma 4.2 which happens with probability at least $1-(5T+3)(M+G+1)\beta^4$. Let us define T_{exh} by $T-M\text{Res}(\beta) = T-7MG\ln(1/\beta)$. Later we will show that this is a lower bound on the time when Conservative Bid Price Control exhausts all free capacity. Define event $\mathcal{S}_1$ by $\{N(g,T) - N(g,T_{\text{exh}}) = 0\}$. That is, there is no group g arrival among the last $7MG\ln(1/\beta)$ arrivals. Then by union bound $\mathbb{P}\{\mathcal{S}_1\} \geq 1 - 7MGp_g\ln(1/\beta)$. In addition, let us define event $\mathcal{S}_g$ (the same as in the proof of Lemma 4.2) by the intersection of the event that $N(g,T) \geq p_gT - 3\sqrt{p_gT\ln(1/\beta)}$ and the event that $N(g,T) - N(g,T-\Delta^{\text{WCB}}) \leq p_g\Delta^{\text{WCB}} + 5\sqrt{\max(p_g\Delta^{\text{WCB}}, \ln(1/\beta))\ln(1/\beta)}$. Using Lemma E.2 and union bound, we have $\mathbb{P}\{\mathcal{S}_g\} \geq 1 - 2\beta^4$. Therefore, $\mathbb{P}\{\mathcal{S}_{\text{BP}} \cap \mathcal{S}_1 \cap \mathcal{S}_g\} \geq 1 - (5T+3)(M+G+1)\beta^4 - 2\beta^4 - 6MGp_g\ln(1/\beta) \geq 1 - (5T+5)(M+G)\beta^4 - 6MGp_g\ln(1/\beta)$. Condition on $\mathcal{S}_{\text{BP}} \cap \mathcal{S}_1 \cap \mathcal{S}_g$. It remains to show that $\alpha_g \geq \min(\mathbb{E}[O_g], O_g) - \frac{12\Delta^{\text{WCB}}}{T}$. If a case of group g in the first T_{emp} periods receives an assignment J^{WAB}, by Lemma 4.2, we have $\alpha_g \geq \mathbb{E}[O_g] - \frac{12\Delta^{\text{WCB}}}{T}$, proving the desired result.

We thus only need to consider the scenario that all cases of group g in the first T_{emp} periods receive their greedy assignments. Without loss of generality, let us assume $N(g,T) \geq 1$ since otherwise $\alpha_g = O_g = 0$. Note that by the ex-post feasibility of the fairness rule (see Definition 1), we must have that O_g is upper bounded by the average greedy score given by $\frac{1}{N(g,T)}\sum_{t\in\mathcal{A}(g,T)} w_{t,J^{\text{GR}}(t)}$. Given the condition that all arrivals of group g in the first T_{emp} periods receive greedy assignments under Conservative Bid Price Control, we have

$$O_g - \alpha_g \leq \frac{\sum_{t\in\mathcal{A}(g,T)} w_{t,J^{\text{GR}}(t)} - \sum_{t\in\mathcal{A}(g,T)} w_{t,J^{\text{CBP}}(t)}}{N(g,T)}.$$

The right hand side is indeed equal to $\frac{\sum_{t\in\mathcal{A}(g,T_{\text{exh}})\setminus\mathcal{A}(g,T_{\text{emp}})}\left(w_{t,J^{\text{GR}}(t)} - w_{t,J^{\text{CBP}}(t)}\right)}{N(g,T)}$ because all cases of group g in the first T_{emp} periods receive their greedy assignments and there is no arrival of group g after period T_{exh} conditioned on $\mathcal{S}_1$. Note that under $\mathcal{S}_{\text{BP}}$, we know $T_{\text{emp}} \geq T - \Delta^{\text{WCB}}$ and thus

$$\sum_{t\in\mathcal{A}(g,T_{\text{exh}})\setminus\mathcal{A}(g,T_{\text{emp}})} \left(w_{t,J^{\text{GR}}(t)} - w_{t,J^{\text{CBP}}(t)}\right) \leq \sum_{t\in\mathcal{A}(g,T_{\text{exh}})\setminus\mathcal{A}(g,T-\Delta^{\text{WCB}})} \left(w_{t,J^{\text{GR}}(t)} - w_{t,J^{\text{CBP}}(t)}\right).$$

For ease of notation, let us define $\mathcal{A}_1 = \mathcal{A}(g,T_{\text{exh}}) \setminus \mathcal{A}(g,T-\Delta^{\text{WCB}})$ by the set of arrivals of group g arriving between $T - \Delta^{\text{WCB}}$ and T_{exh}. It remains to upper bound $\frac{\sum_{t\in\mathcal{A}_1}\left(w_{t,J^{\text{GR}}(t)} - w_{t,J^{\text{CBP}}(t)}\right)}{N(g,T)}$.

Consider two scenarios. First, if $p_g\Delta^{\text{WCB}} \geq \ln(1/\beta)$, we have $p_gT \geq 36p_g\Delta^{\text{WCB}}$ since $T \geq 36\Delta^{\text{WCB}}$ by assumption. It also implies $p_gT \geq 6\sqrt{p_gT\ln(1/\beta)}$. By the definition of $\mathcal{S}_g$, we know $N(g,T) \geq p_gT - 3\sqrt{p_gT\ln(1/\beta)} \geq p_gT - \frac{1}{2}p_gT = \frac{1}{2}p_gT$. In addition, the second event in $\mathcal{S}_g$ together with $p_g\Delta^{\text{WCB}} \geq \ln(1/\beta)$ gives $N(g,T) - N(g,T-\Delta^{\text{WCB}}) \leq p_g\Delta^{\text{WCB}} + 5\sqrt{p_g\Delta^{\text{WCB}}\ln(1/\beta)} \leq 6p_g\Delta^{\text{WCB}}$. Therefore, assuming that $p_g\Delta^{\text{WCB}} \geq \ln(1/\beta)$, we have

$$O_g - \alpha_g \leq \frac{\sum_{t\in\mathcal{A}_1}\left(w_{t,J^{\text{GR}}(t)} - w_{t,J^{\text{CBP}}(t)}\right)}{N(g,T)} \leq \frac{N(g,T) - N(g,T-\Delta^{\text{WCB}})}{N(g,T)} \leq \frac{6p_g\Delta^{\text{WCB}}}{0.5p_gT} = \frac{12\Delta^{\text{WCB}}}{T}.$$

A more difficult scenario is $p_g\Delta^{\text{WCB}} \leq \ln(1/\beta)$. That is, the group is a small group. We make the following claim that under $\mathcal{S}_{\text{BP}} \cap \mathcal{S}_1 \cap \mathcal{S}_g$, all group g arrivals receive assignments J^{WCB} and never face the scenario of no capacity.

Claim 1. *Conditioned on $\mathcal{S}_{\text{BP}} \cap \mathcal{S}_1 \cap \mathcal{S}_g$, we have $J^{\text{CBP}}(t) = J^{\text{WCB}}(t)$ for every $t \in \mathcal{A}(g,T)$.*

Under this claim, we have that either all arrivals of group g receive greedy assignment; or for at least one arrival t, condition (predict-to-meet) holds and $J^{\text{WCB}}(t) = J^{\text{WAB}}(t)$. In the latter scenario, using Lemma 4.2 gives $\alpha_g \ge \mathbb{E}[O_g] - \frac{12\Delta^{\text{WCB}}}{T}$. For the first scenario, we have $\alpha_g = \frac{1}{N(g,T)} \sum_{t \in \mathcal{A}(g,T)} w_{t,J^{\text{GR}}(t)} \ge O_g$. Summarizing the above, we then have for a group g with $p_g \Delta^{\text{WCB}} \le \ln(1/\beta)$, we always have $\alpha_g \ge \min\left(\mathbb{E}[O_g] - \frac{12\Delta^{\text{WCB}}}{T}, O_g\right)$.

It remains to prove the claim. By $\mathcal{S}_g$, we have $N(g,T) - N(g, T - \Delta^{\text{WCB}}) \le \ln(1/\beta) + 5\ln(1/\beta) = 6\ln(1/\beta)$. Since $T_{\text{emp}} \ge T - \Delta^{\text{WCB}}$ under $\mathcal{S}_{\text{BP}}$, we also have $N(g,T) - N(g, T_{\text{emp}}) \le 6\ln(1/\beta)$. Therefore, the number of group g arrivals is at most $6\ln(1/\beta)$ after the time one location runs out of free capacity. Recall that there is no group g arrivals during $[T_{\text{exh}}+1, T]$. Denote by n the number of group g cases during $[T_{\text{emp}}+1, T]$ and their labels by $T_{\text{emp}} < t_1 < \ldots < t_n \le T_{\text{exh}}$. We next show that for all these arrivals, there is remaining reserved capacity. That is, $r_{g,j}(t_k) \ge 1$ for every $j \in \mathcal{M}, k \in [n]$. To show this, by Conservative Bid Price Control, we only use reserves after T_{emp} since the algorithm first consumes free capacity. In addition, the algorithm does not recollect reserved capacity until there is no free capacity for any location (see Line 10 in Algorithm 2). We know the amount of free capacity at the beginning is $\sum_{j \in \mathcal{M}} f_j(1) \ge \sum_{j \in \mathcal{M}} (s_j - 7G\ln(1/\beta)) = T - 7MG\ln(1/\beta)$ and each case decreases the capacity by at most one. Therefore, for $t \le T_{\text{exh}} = T - 7MG\ln(1/\beta)$, we must have $\sum_{j \in \mathcal{M}} f_j(t) \ge 1$. It then implies that for $k \in [n]$, $c_g(t_k) \ge \text{Cap}(\beta) - (k-1) \ge 6\ln(1/\beta) - (n-1) \ge 1$, where the second inequality is by the selection of reserve cap in (PAR). In addition, by the setting of reserve $\text{Res}(\beta)$ in a location in (PAR), the number of reserves is enough for every group to meet their cap. As a result, a group g arrival never sees the lack of reserve in a location and the limit of the reserve cap, i.e., Line 7 in Algorithm 2 is invoked throughout for every greedy steps of group g. We then have that if case t is of group g, Conservative Bid Price Control assigns $J^{\text{CBP}}(t) = J^{\text{WCB}}(t)$ for $t \in \mathcal{A}(g,T)$.

Putting together the proof, we know that condition on $\mathcal{S}_{\text{BP}} \cap \mathcal{S}_1 \cap \mathcal{S}_g$, for a group g, there are three scenarios. First, an arrival before T_{emp} receives an assignment J^{WAB} and then $\alpha_g \ge \mathbb{E}[O_g] - \frac{12\Delta^{\text{WCB}}}{T}$. Second, all arrivals before T_{emp} receive greedy assignments and $p_g\Delta^{\text{WCB}} \ge \ln(1/\beta)$. Under this scenario, we show $\alpha_g \ge O_g - \frac{12\Delta^{\text{WCB}}}{T}$. Finally, if $p_g\Delta^{\text{WCB}} \le \ln(1/\beta)$, we show $\alpha_g \ge \min\left(\mathbb{E}[O_g] - \frac{12\Delta^{\text{WCB}}}{T}, O_g\right)$. As a result, for a group g, we have $\alpha_g \ge \min(\mathbb{E}[O_g], O_g) - \frac{12\Delta^{\text{WCB}}}{T}$ with probability lower bounded by $\mathbb{P}\{\mathcal{S}_{\text{BP}} \cap \mathcal{S}_1 \cap \mathcal{S}_g\} \ge 1 - (5T+5)(M+G)\beta^4 - 7MGp_g\ln(1/\beta)$. $\square$

E Concentration inequalities

In the proof, we frequently use Hoeffding's Inequality simplified from [BLM13].

Fact 3 (Hoeffding's Inequality)**.** *Given n independent random variables X_i taking values in $[0,1]$ almost surely. Let $X = \sum_{i=1}^n X_i$. Then for any $x > 0$,*

$$\mathbb{P}\{X < \mathbb{E}[X] - x\} \le e^{-2x^2/n};\ \mathbb{P}\{X > \mathbb{E}[X] + x\} \le e^{-2x^2/n}. \tag{20}$$

Another useful inequality is the Bernstein's Inequality. The following version is by Corollary 2.11 from [BLM13].

Fact 4 (Bernstein's Inequality)**.** *Given n independent random variables X_i such that $X_i \le b$ almost surely and $\sum_{i=1}^n \mathbb{E}\left[X_i^2\right] \le \nu$. Let $X = \sum_{i=1}^n X_i$. Then for all $x > 0$,*

$$\mathbb{P}\{X > \mathbb{E}[X] + x\} \le \exp\left(\frac{-x^2}{2(\nu + bx/3)}\right).$$

An implication of the above result is the following lemma.

Lemma E.1. *Given n independent random variables X_i such that $X_i \geq 0$ almost surely and $\sum_{i=1}^n \mathbb{E}\left[X_i^2\right] \leq \nu$. Let $X = \sum_{i=1}^n X_i$. Then for all $x > 0$,*

$$\mathbb{P}\{X < \mathbb{E}\left[X\right] - x\} \leq \exp\left(\frac{-x^2}{2\nu}\right).$$

Proof. Take $Y_i = -X_i$. Then $\sum_{i=1}^n \mathbb{E}\left[Y_i^2\right] \leq \nu$ and $Y_i \leq 0$ almost surely. By Fact 4, for any $x > 0$, we have $\mathbb{P}\{\sum_{i=1}^n Y_i > \mathbb{E}\left[\sum_{i=1}^n Y_i\right] + x\} \leq \exp\left(\frac{-x^2}{2\nu}\right)$. Rearranging terms give the desired result. $\square$

We finally recall the Chernoff bound of Bernoulli random variables and an implied version of it. We summarize the result as follows.

Lemma E.2. *Given a binomial random variable X, for any $x > 0$, we have*

$$\mathbb{P}\{X < \mathbb{E}\left[X\right] - x\} \leq \exp\left(-x^2/(2\mathbb{E}\left[X\right])\right); \ \mathbb{P}\{X > \mathbb{E}\left[X\right] + x\} \leq \exp\left(-\frac{x^2}{2(\mathbb{E}\left[X\right] + x/3)}\right).$$

In particular, the second inequality implies that for any $\beta > 0$, we have

$$\mathbb{P}\left\{X > \mathbb{E}\left[X\right] + 5\sqrt{\max(\mathbb{E}\left[X\right], \ln\frac{1}{\beta})\ln\frac{1}{\beta}}\right\} \leq \beta^4.$$

Proof. The first two probability bounds are implied by the Bernstein's Inequality (Fact 4) since X is a sum of Bernoulli random variables. We prove the third bound here. By the second probability bound, we have

$$\begin{aligned}
&\mathbb{P}\left\{X > \mathbb{E}\left[X\right] + 5\sqrt{\max(\mathbb{E}\left[X\right], \ln\frac{1}{\beta})\ln\frac{1}{\beta}}\right\} \\
&\leq \exp\left(-\frac{25\max(\mathbb{E}\left[X\right], \ln\frac{1}{\beta})\ln\frac{1}{\beta}}{2\mathbb{E}\left[X\right] + 4\sqrt{\max(\mathbb{E}\left[X\right], \ln\frac{1}{\beta})\ln\frac{1}{\beta}}}\right) \\
&\leq \exp\left(-\frac{25\max(\mathbb{E}\left[X\right], \ln\frac{1}{\beta})\ln\frac{1}{\beta}}{6\max(\mathbb{E}\left[X\right], \ln\frac{1}{\beta})}\right) \leq \exp\left(-4\ln\frac{1}{\beta}\right) = \beta^4.
\end{aligned}$$

$\square$