Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Sequential Fair Allocation: Achieving the Optimal Envy-Efficiency Trade-off Curve

Sean R. Sinclair, Gauri Jain, Siddhartha Banerjee, Christina Lee Yu

informs
https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Articles in Advance, pp. 1–17
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Crosscutting Areas

Sequential Fair Allocation: Achieving the Optimal Envy-Efficiency Trade-off Curve

Sean R. Sinclair,[a,*] Gauri Jain,[a] Siddhartha Banerjee,[a] Christina Lee Yu[a]

[a] Operations Research and Information Engineering, Cornell University, Ithaca, New York 14853

*Corresponding author

Contact: srs429@cornell.edu, https://orcid.org/0000-0002-7011-8253 (SRS); gj82@cornell.edu, https://orcid.org/0000-0002-7517-3593 (GJ); sbanerjee@cornell.edu, https://orcid.org/0000-0002-8954-4578 (SB); cleeyu@cornell.edu, https://orcid.org/0000-0002-2165-5220 (CLY)

Received: May 11, 2021
Revised: October 11, 2021; December 20, 2021; July 19, 2022
Accepted: September 29, 2022
Published Online in Articles in Advance: November 23, 2022

Area of Review: Revenue Management and Market Analytics

https://doi.org/10.1287/opre.2022.2397

Abstract. We consider the problem of dividing limited resources to individuals arriving over T rounds. Each round has a random number of individuals arrive, and individuals can be characterized by their type (i.e., preferences over the different resources). A standard notion of fairness in this setting is that an allocation simultaneously satisfy envy-freeness and efficiency. The former is an individual guarantee, requiring that each agent prefers the agent's own allocation over the allocation of any other; in contrast, efficiency is a global property, requiring that the allocations clear the available resources. For divisible resources, when the number of individuals of each type are known up front, the desiderata are simultaneously achievable for a large class of utility functions. However, in an online setting when the number of individuals of each type are only revealed round by round, no policy can guarantee these desiderata simultaneously, and hence, the best one can do is to try and allocate so as to approximately satisfy the two properties. We show that, in the online setting, the two desired properties (envy-freeness and efficiency) are in direct contention in that any algorithm achieving additive counterfactual envy-freeness up to a factor of L_T necessarily suffers an efficiency loss of at least $1/L_T$. We complement this uncertainty principle with a simple algorithm, Guarded-Hope, which allocates resources based on an adaptive threshold policy and is able to achieve any fairness–efficiency point on this frontier. Our results provide guarantees for fair online resource allocation with high probability for multiple resource and multiple type settings. In simulation results, our algorithm provides allocations close to the optimal fair solution in hindsight, motivating its use in practical applications as the algorithm is able to adapt to any desired fairness efficiency trade-off.

Funding: This work was supported by the National Science Foundation [Grants ECCS-1847393, DMS-1839346, CCF-1948256, and CNS-1955997] and the Army Research Laboratory [Grant W911NF-17-1-0094].
Supplemental Material: The online appendix is available at https://doi.org/10.1287/opre.2022.2397.

Keywords: online resource allocation • Varian fairness • Nash social welfare • model-predictive control

1. Introduction

Our work here is motivated by a problem faced by a collaborating food bank (Food Bank for the Southern Tier of New York (FBST; https://www.foodbankst.org/) in operating its mobile food pantry program. Recent demands for food assistance have climbed at an enormous rate, and an estimated 14 million children are not getting enough food because of the COVID-19 epidemic in the United States (Bauer 2020, Kulish 2020). With sanctions on operating in-person stores, many food banks have increased their mobile food pantry services. In these systems, the mobile food pantry must decide on how much food to allocate to a distribution center on arrival without knowledge of demands in future locations. This model also extends as a representation of broader stockpile allocation problems (such as vaccine and medical supply allocation) and reservation mechanisms.

As a simplified example (see Section 3 for the full model, including multiple resources and individual types), every day, the mobile food pantry uses a truck to deliver B units of food supplies to individuals over T rounds (each round can be thought of as a distribution location: soup kitchens, pantries, nursing homes, etc.). When the truck arrives at a site t (or round t), the operator observes N_t individuals and chooses how much to allocate to each individual ($X_t \in \mathbb{R}^{N_t}$) before moving to the next round. The number of people assembling at each site changes from day to day, and the operator typically does not know the number of individuals at later sites (but has a sense of the distribution based on previous visits).

In off-line problems, in which the number of individuals at each round $(N_t)_{t\in[T]}$ is known to the principal in advance, there are many well-studied notions of fair allocation of resources. One guarantee, envy-freeness, requires that each individual prefers the individual's own allocation over the allocation of any other. In contrast, efficiency is a global property, requiring that the allocations clear the available resources. For divisible resources, these desiderata are simultaneously achievable for a large class of utility functions with multiple resources and easily computed (via a convex program) by maximizing the Nash social welfare (NSW) objective subject to allocation constraints (Eisenberg 1961, Varian 1974). As an example, in this (simplified) setting, the fair allocation is easily computed by allocating $X^{opt} = B/N$ to each individual, where $N = \sum_{t\in[T]} N_t$ is the total number of individuals across all rounds. This allocation is clearly envy-free (as each individual receives an equal allocation) and is efficient (as all of the resources are exhausted); it's also easy to see that this is the only allocation that satisfies these two properties simultaneously.

Many practical settings, however, operate more akin to the FBST mobile food pantry, in which the principal makes allocation decisions online with incomplete knowledge of the demand for future locations. However, these principals do have access to historical data allowing them to generate histograms over the number of individuals for each round (or potentially just first moment information). Designing good allocation algorithms in such settings necessitates harnessing the Bayesian information of future demands to ensure equitable access to the resource and also adapting to the online realization of demands as they unfold to ensure efficiency.

Satisfying any one of these properties is trivially achievable in online settings. The solution that allocates $X_t = 0$ to each individual satisfies hindsight envy-freeness as each individual is given an equal allocation. The solution that allocates $X_1 = B/N_1$ to individuals at the first location and $X_t = 0$ for $t \geq 2$ satisfies efficiency as the entire budget is exhausted at the first location. Another difficult challenge in this setting is achieving low counterfactual envy, ensuring that the allocations made by the algorithm (X_t) are close to what each individual should have received with the fair solution in hindsight (B/N). More meaningful is understanding how these different criteria interact. Here, we tackle these important challenges by defining meaningful notions of approximately fair online allocations and develop algorithms that are able to utilize distributional knowledge to achieve allocations that strike a balance between the competing objectives of envy and efficiency.

1.1. Overview of Our Contributions

In sequential settings, one way to measure the (un)fairness of any online allocation (X^{alg}) is in terms of its counterfactual distance (for both envy and efficiency) when compared with the optimal fair allocation in hindsight (i.e., off-line allocation X^{opt}). Another measure is hindsight envy (when compared only to allocations made by the algorithm). In particular, we define *counterfactual envy* as $\Delta_{EF} = \|u(X_\theta^{opt}, \theta) - u(X_\theta^{alg}, \theta)\|_\infty$ to be the maximum difference in utility between the algorithm's allocation and the off-line allocation when agents are characterized by their type θ and define *hindsight envy* as $\text{ENVY} = \max_{t,t',\theta,\theta'} u(X_{t',\theta'}^{alg}, \theta) - u(X_{t,\theta}^{alg}, \theta)$ to be the maximum difference between the utility individuals would receive if given someone else's allocations and let $\Delta_{efficiency} = B - \sum_{t,\theta} N_{t,\theta} X_{t,\theta}^{alg}$ be the algorithm's total leftover resources. These are all very stringent metrics, akin to the notion of regret in online decision-making settings.

In these settings with competing objectives, practitioners often resort to ad hoc rules of thumb, heuristics, and trial-and-error adjustments of the system to attempt to manage the balance between objectives. How these criteria interact and trade off among one another is often not well-understood or characterized, and furthermore, there typically does not exist a single best "ranking" or a clear single objective function that determines which trade-offs are better than others. In fact, minimizing some combination of $(\Delta_{EF}, \text{ENVY}, \Delta_{efficiency})$ can be formulated as a Markov decision process (MDP). However, as these metrics depend on the entire allocation, the complexity of finding the optimal policy is exponential in the number of rounds and may be difficult to interpret (Manshadi et al. 2021). Moreover, it is much harder to use MDP formulations to explore the trade-off between the objectives.

Our main technical contribution is to provide a complete characterization of the achievable pairs of $(\Delta_{EF}, \text{ENVY}, \Delta_{efficiency})$. Our results hold in expectation and with high probability under multiple divisible resources and with a finite set of individual types with linear utilities. In particular, we show the following informal theorem (see Figure 1 for a graphic representation).

Informal Theorem 1 (See Sections 4 and 7 for Full Versions). *Under mild regularity conditions on the distribution of N_t, we have the following ($\gtrsim$ ignores problem-dependent constants, logarithmic factors of T, and $o(1)$ factors):*

1. *(Statistical uncertainty principle) Any online allocation algorithm must suffer counterfactual envy of at least $\Delta_{EF} \gtrsim 1/\sqrt{T}$.*

2a. *(Counterfactual envy-efficiency uncertainty principle) Any online allocation algorithm necessarily suffers $\Delta_{efficiency} \gtrsim \min\{\sqrt{T}, 1/\Delta_{EF}\}$.*

2b. *(Hindsight envy-efficiency uncertainty principle) Any online allocation algorithm necessarily suffers $\Delta_{efficiency} \gtrsim \min\{\sqrt{T}, 1/\text{ENVY}\}$.*

3. *(Upper Bound via GUARDED-HOPE) For any choice of L_T, with probability at least $1-\delta$, GUARDED-HOPE with parameter L_T achieves*

$$\text{ENVY} \leq L_T \qquad \Delta_{EF} \lesssim \max\{1/\sqrt{T}, L_T\}$$

$$\Delta_{efficiency} \lesssim \min\{\sqrt{T}, 1/L_T\}.$$

Figure 1. (Color online) Graphic Representation of the Major Contributions (Theorem 1)

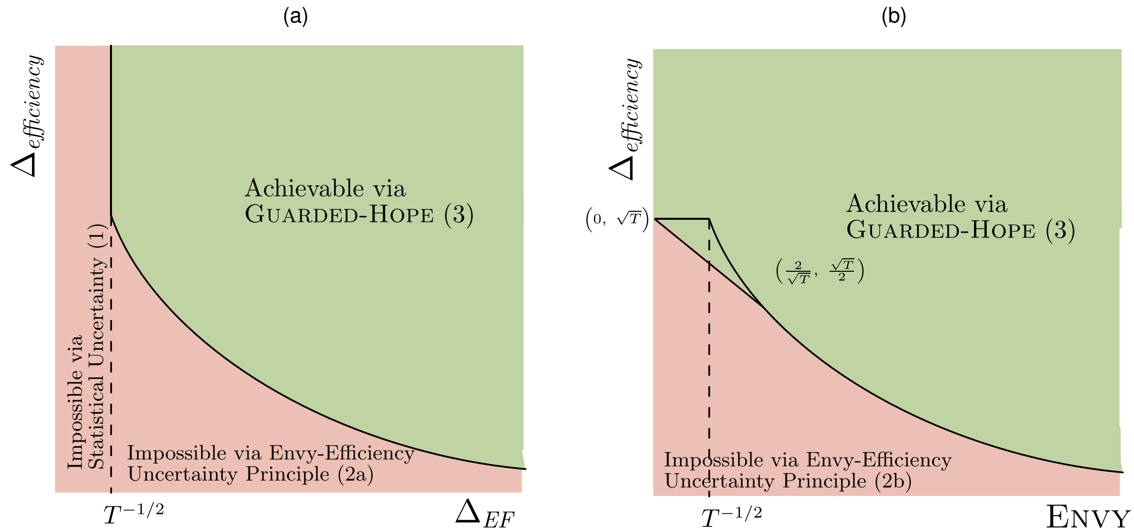

Notes. Here, the x-axis denotes Δ_{EF} or ENVY, and the y-axis denotes $\Delta_{efficiency}$, the remaining resources. The dotted line represents the impossibility resulting from statistical uncertainty in the optimal allocation, and the region below the solid line represents the impossibility resulting from the envy-efficiency uncertainty principle. (a) $\Delta_{EF} - \Delta_{efficiency}$. (b) ENVY $- \Delta_{efficiency}$.

In short, our results show that envy and waste must be inversely proportional to one another such that decreasing envy requires increasing waste and vice versa. The lower bounds (1 and 2) are established using anti-concentration arguments alongside understanding the fundamental gap in ensuring enough resources to allocate close to the estimated optimal solution while simultaneously trying to eliminate waste.

Furthermore, we provide a simple algorithm, GUARDED-HOPE, which achieves the correct trade-off between envy and waste, matching the lower bound in terms of T up to logarithmic factors. Given an input of L_T, our algorithm satisfies a hindsight envy bound of ENVY $\lesssim L_T$ and counterfactual envy bound of $\Delta_{EF} \lesssim \max\{1/\sqrt{T}, L_T\}$ with waste bounded by $\Delta_{efficiency} \lesssim \max\{\sqrt{T}, 1/L_T\}$. Our algorithm achieves this using novel concentration arguments on the optimal Nash social welfare solution, utilizing a sensitivity argument on the solution to the optimization problem instead of the objective (as commonly used for competitive ratio guarantees) to learn a lower guardrail on the optimal solution in hindsight. Given this, we construct an upper guardrail to satisfy the desired Δ_{EF} and ENVY bound. We then achieve the proper trade-off by carefully balancing allocating the established lower guardrail with the upper guardrail and simultaneously ensuring the algorithm never runs out of budget.

To get some intuition into the envy-efficiency uncertainty principle, consider the simple food bank example described for a single resource (with arrivals N_t in each location and $X^{opt} = B/N$, where $N = \sum_{t\in[T]} N_t$). For convenience, we temporarily assume that each agent's utility is directly proportional to the agent's allocation (i.e., $u(X, \theta) = X$). Consider allocation X_1 at the first location: via standard concentration arguments, one can find a high probability lower confidence bound for B/N with a half-width on the order of $1/\sqrt{T}$. Now, it's not hard to argue that allocating according to the lower confidence bound at all locations achieves counterfactual envy of $\Delta_{EF} \approx 1/\sqrt{T}$, ENVY $= 0$, and $\Delta_{efficiency} \approx \sqrt{T}$. This corresponds to the cusp of the efficiency-envy trade-off curves in Figure 1.

Now, if we relax the Δ_{EF} or ENVY constraint to $\approx 1/T^{1/3}$ and use the naive static policy of always allocating via the now looser lower confidence bound, we get a waste of $T \cdot T^{-1/3} = T^{2/3}$. Our algorithm instead takes a different approach, using the lower confidence bound of order $1/\sqrt{T}$ as the lower guardrail allocation, and sets the upper guardrail allocation to be the lower one plus the desired bound on Δ_{EF} or ENVY. If we establish that the algorithm always allocates within the guardrails, we automatically have the desired bound on Δ_{EF} and ENVY. The main additional factor in achieving the trade-off for $\Delta_{efficiency}$ is ensuring we properly allocate according to the upper threshold and ensure we do not run out of budget to ensure the lower threshold allocation. With this, GUARDED-HOPE achieves $\Delta_{efficiency} \approx T^{1/3}$, which furthermore is the best possible. Moreover, we complement our theoretical results with experiments highlighting the empirical performance of different algorithms (on both synthetic settings as well as a data set based on mobile food pantry operations), which shows that GUARDED-HOPE has much lower waste and envy compared with static under-allocation as well as other certainty equivalence based heuristics Bertsekas (2012).

Whereas fairness in resource allocation is well-studied in off-line and adversarial settings, fairness metrics for the sequential stochastic setting are poorly understood (especially when individuals are arriving online). Our proposed metrics and results give a novel way of extending Varian's definitions of fairness to the sequential setting. Moreover, ours is the first result to provide guarantees for fair online resource allocation with high probability for multiple-resource and multiple-type settings. Most existing work aims to show competitive ratio or additive guarantees on the Nash social welfare objective (Banerjee et al. 2020) or focus on the max-min objective (Lien et al. 2014, Manshadi et al. 2021). Such guarantees are dangerously misleading in that the resultant allocations may exhibit clear unfairness in hindsight. Similarly, an ex ante or probabilistic guarantee may also be perceived as unfair; both allocating one unit with certainty and allocating 10 units with probability 1/10 give the same ex ante guarantee. In contrast, our chosen metrics and theoretical results provides a firm basis for counterfactual and ex post individual fairness guarantees. Whereas we do not believe our work gives a final answer in the theoretical and practical understanding of fairness in online allocation, we hope it adds to the conversation of incorporating ethics into sequential AI algorithms. More discussion on the advantages and disadvantages of our proposed fairness metrics is in Online Appendix A.

1.2. Other Motivating Examples

In addition to the mobile food pantry allocation problem that forms the focus of our work, we believe our ideas can prove useful in several other settings:

1.2.1. Stockpile Allocation. In many healthcare systems or resource allocation problems, government mechanisms decide how to allocate critical resources to states, individuals, or hospitals. For example, the U.S. federal government was tasked with distributing Remdesivir, an antiviral drug used early in the pandemic for COVID-19 treatment (Lupkin 2020). More recently relevant, states and government organizations are deciding how to allocate COVID-19 (or influenza) vaccines to various population demographics across several rounds (Jaberi-Douraki and Moghadas 2014, Yi and Marathe 2015). In these scenarios, on a monthly basis, each state is given a fixed amount of the resource (say COVID-19 vaccinations) and is tasked with distributing these to individuals across various distribution locations. Whereas the primary goal is to develop efficient allocations, an alternative objective may be to ensure equitable access to the resource (Donahue and Kleinberg 2020, Shadmi et al. 2020, Manshadi et al. 2021).

1.2.2. Reservation Mechanisms. These are key for operating shared high-performance computing (HPC) systems (Ghodsi et al. 2011). Cluster centers for HPC receive numerous requests online with varying demands for CPUs and graphics processing units (GPUs). Algorithms must allocate resources to incoming jobs with only distributional knowledge of future resource demands. Important to these settings is the large number of resources (number of GPUs, RAM, etc., available at the center), requiring algorithms that scale to higher dimensional problems.

2. Related Work

Fairness in resource allocation and the use of Nash social welfare was pioneered by Varian in his seminal works (Varian 1974, 1976). Since then, researchers have investigated fairness properties for both off-line and online allocation in settings with divisible or indivisible resources and when either the individuals or resources arrive online. We now briefly discuss some related works; see Aleksandrov and Walsh (2020) for a comprehensive survey. What distinguishes our setting from many of the previous works is that we consider the online Bayesian setting with a known distribution. Many previous works are either limited to off-line or nonadaptive algorithms or consider adversarial online arrivals. Trade-offs between various fairness metrics is also considered previously in the literature but for classification-based fairness metrics on protected attributes instead of allocation-based ones (Kleinberg et al. 2016).

2.1. Food Bank and Healthcare Operations

There is a growing body of work in the operations research literature addressing logistics and supply chain issues in the area of humanitarian relief, healthcare, and food distribution (Jaberi-Douraki and Moghadas 2014, Yi and Marathe 2015, Orgut et al. 2016, Sengul Orgut et al. 2017, Alkaabneh et al. 2020). The research focuses on designing systems that balance efficiency, effectiveness, and equity. In Eisenhandler and Tzur (2019), they study the logistical challenges of managing vehicles with limited capacity to distribute food and provide routing and scheduling protocols. In Lien et al. (2014) and Manshadi et al. (2021), they consider sequential allocation with an alternative objective of maximizing the minimum utility (also called the leximin in the literature; Moulin 2004). We instead consider sequential allocation of resources under the objectives of achieving approximate fairness notions with regards to envy and efficiency.

2.2. Cake Cutting

Cake cutting serves as a model for dividing a continuous object (whether that be a cake, advertisement space, land, etc.) (Brams and Taylor 1995, Procaccia 2013). Under this model, prior work considers situations in which individuals arrive and depart during the process of dividing a resource, and the utility of an agent is a set function on the interval of the resource received. Researchers analyze the off-line setting to develop algorithms to allocate the resource with a minimal number of cuts (Brams and Taylor 1996) or online under adversarial arrivals (Walsh 2011). Our model instead imposes stochastic assumptions on the number of arriving individuals and characterizes probabilistic instead of sample-path fairness criteria.

2.3. Online Resources

One line of work considers the resource (here, to be thought of as the units of food, processing power, etc.) is online and the agents are fixed (Aleksandrov et al. 2015; Mattei et al. 2017, 2018; Benade et al. 2018; Aleksandrov and Walsh 2019; Banerjee et al. 2020; Bansal et al. 2020; Bogomolnaia et al. 2022). In Zeng and Psomas (2020), they study the trade-offs between fairness and efficiency when items arrive under several adversarial models. Another common criterion is designing algorithms that are envy-free up to one item, for which researchers design algorithms that can reallocate previously allocated items but try to minimize these adjustments (Aziz et al. 2016, He et al. 2019). These problems are in contrast to our model in which, instead, the resources are fixed and depleting over time and individuals arrive online.

2.4. Online Individuals

The other setting more similar to our work considers agents as arriving online and the resources as fixed. In Kalinowski et al. (2013), they consider this setting in which the resources are indivisible with the goal of maximizing utilitarian welfare (or the sum of utilities), which provides no guarantees on individual fairness. Another approach in Gerding et al. (2019) considers a scheduling setting in which agents arrive and depart online. Each agent has a fixed and known arrival time, departure time, and demand. The goal then is to determine a schedule and allocation that is Pareto-efficient and envy-free. Another line of work (Cole et al. 2013; Friedman et al. 2015, 2017) considers fair division with minimal disruptions on previous allocations. Their fairness ratio can be viewed as a competitive ratio form of our counterfactual envy definition (Definition 2).

2.5. Nonadaptive Allocations

A separate line of research considers fairness questions for resource allocations in a similar setting in which the utilities across groups are drawn from known probability distributions (Elzayn et al. 2019, Donahue and Kleinberg 2020). They investigate probabilistic versions of fairness, in which the goal is to quantify the discrepancy between the objectives of ensuring the expected utilization of the resources is large (ex ante Pareto-optimal), whereas the probability of receiving the resource is proportional across groups (ex ante proportional). However, they consider algorithms that decide on the entire allocation for each agent up front before observing the demand rather than adaptive policies.

2.6. Adaptive Allocations

In contrast, we consider a model in which the principal makes decisions on the amount of resources after witnessing the number of individuals in a round. Most similar to our work is recent work analyzing a setting in which individuals arrive over time and do not depart so that the algorithm can allocate additional resources to individuals who arrived in the past (Kash et al. 2014). We instead consider a stochastic setting in which individuals arrive and depart in the same step with the goal of characterizing allocations that cannot reallocate to previous agents. Other papers either seek competitive ratios in terms of the Nash social welfare objective (Azar et al. 2010, Banerjee et al. 2020, Bateni et al. 2022) or derive allocation algorithms that perform well in terms of max-min (Lien et al. 2014, Manshadi et al. 2021). Our work differs from these in that we impose additional distribution assumptions (notably that the variance of the demand is on a smaller order than its mean, more common in real-world scenarios). The results in Manshadi et al. (2021) can be viewed as highlighting a trade-off between efficiency and the max-min objective although achieving efficiency of zero is trivial in that setting as the algorithm designer is not penalized for giving all leftover resources at the last location. In contrast, under our setting, eliminating the resources at the final round penalizes the algorithm in terms of both Δ_{EF} and ENVY, requiring a more nuanced discussion on the trade-off between efficiency and envy.

3. Preliminaries

We use $\mathbb{R}_+$ to denote the set of nonnegative reals, $\|X\|_\infty = \max_{i,j}|X_{i,j}|$ to denote the matrix maximum norm, and cX to denote entry-wise multiplication for a constant c. When comparing vectors, we use $X \le Y$ to denote that each component $X_i \le Y_i$.

3.1. Model and Assumptions

A principal is tasked with dividing K divisible resources among a population of individuals who are divided between T distinct rounds; these can represent T locations visited sequentially by the principal (for example, food distribution sites visited by a mobile pantry) or T consecutive time periods (for example, days over which a hospital must stretch some limited medical supply before it is restocked).

Each resource $k \in [K]$ has a fixed initial budget B_k that the principal can allocate across these rounds. Each round has a (possibly random) set of distinct individuals arriving to request a share of the resources. Individuals are characterized by their type $\theta \in \Theta$, corresponding to their preferences over the K resources, in which individuals of type θ receive utility $u(x, \theta) : \mathbb{R}^K \times \Theta \to \mathbb{R}$ for an allocation x. We henceforth assume that the set of possible types has finite cardinality $|\Theta|$ and denote $(N_{t,\theta})_{\theta\in\Theta}$

to be the vector containing the number of arrivals of each type in round t, the demand $N_{t,\theta}$ denotes the number of type-θ arrivals. $(N_{t,\theta})_{\theta\in\Theta,t\le T}$ is drawn from some known distribution $\mathcal{F}$; note that these distributions across rounds need not be identical.

In the ex post or off-line setting, the number of individuals per round $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$ is known in advance and can be used by the principal to choose allocations $X \in \mathbb{R}^{T\times|\Theta|\times K}$ for individuals in each round t of type θ. In the online setting, the principal considers each round sequentially in a fixed order $t = 1,\dots,T$, is informed of the number of individuals $(N_{t,\theta})_{\theta\in\Theta}$ in that round, and chooses allocation $X_t^{alg} \in \mathbb{R}^{|\Theta|\times K}$ before continuing on to the next round in which $X_{t,\theta,k}^{alg}$ denotes the allocation of resource k earmarked for each of the $N_{t,\theta}$ individuals of type θ in that round. This assumption includes not only independent and identically distributed (i.i.d.) demands, but can also be extended to distributions that arise from Markov chains or latent variable models (see Section 7 for more details). We impose the additional assumption that the algorithm allocates the same allocation to each of the $N_{t,\theta}$ individuals of type θ. This is without loss of generality as one of the primary goals of the paper is to investigate envy, whereby one out of any two individuals of type θ in round t envies the other unless their allocations are the same. Allocation decisions are irreversible and must obey the overall budget constraints.

3.1.1. Assumptions. We assume that, for every $t \in [T]$ and $\theta \in \Theta$, $N_{t,\theta} \ge 1$ almost surely. We also assume that $N_{t,\theta}$ are independent with variance $\mathrm{Var}[N_{t,\theta}] = \sigma_{t,\theta} > 0$ and mean absolute deviation $|N_{t,\theta} - \mathbb{E}[N_{t,\theta}]| = \rho_{t,\theta} < \infty$ almost surely. We additionally denote $\sigma^2_{min} = \min_{t,\theta}\sigma^2_{t,\theta}$, $\sigma^2_{max} = \max_{t,\theta}\sigma^2_{t,\theta}$ and $\mu_{max} = \max_{t,\theta}\mathbb{E}[N_{t,\theta}]$ and assume that $\sigma^2_{min}, \sigma^2_{max}, \mu_{max}$ are given constants. These assumptions are for ease of notation and clarity of presentation; in particular, our results only depend on mild conditions on the expectation and tails of the sums of future arrivals $\sum_{t'>t} N_{t',\theta}$ of each type. Extensions are discussed in Section 7. We define $\beta_{avg} = B/\sum_{\theta\in\Theta}\sum_{t\in[T]}\mathbb{E}[N_{t,\theta}] \in \mathbb{R}^k$ as the average resource per individual; for ease of understanding, β_{avg} can be viewed as being a constant, but our results hold for any β_{avg}.

We also focus on utility functions that are linear, that is, for which $u(x,\theta) = \langle w_\theta, x\rangle$, where the latent individual type θ is characterized by $w_\theta \in \mathbb{R}^K_{\ge 0}$ as a vector of preferences over each of the different resources. For example, the type θ could refer to a "vegetarian" type with preferences $[2,0,1]$ over the set of resources [produce, meat, canned soup] indicating a marginal utility of zero for any allocated meat and increased preference for produce. The relative scale of the weights help indicate preference for one food resource over another.

The assumption that agents' preferences over resources are linear is limiting in that it does not account for settings in which resources exhibit complementarities (modeled via, e.g., Leontief, or filling, utilities) in addition to omitting popular utility functions in the extant literature (e.g., Cobb–Douglas utilities). Our algorithmic techniques naturally extend to more general utility functions (so long as the Eisenberg–Gale (EG) program can be solved efficiently). However, we leave understanding both the upper and lower bounds on the achievable envy and efficiency pairs to future work. More details on modeling individual utilities for the experiments are in Online Appendix C.

Finally, we assume that our resources are divisible, in that allocations can take values in $\mathbb{R}^K_+$. In our particular regimes of interest in which we scale the number of rounds and budgets, this is easy to relax to integer allocations with vanishing loss in performance.

3.1.2. Additional Notation. We use $B = (B_1,\dots,B_K)$ to be the budget vector. For any location t and type θ, we use $N_{\ge t,\theta}$ to denote $\sum_{t'\ge t} N_{t',\theta}$. If the subscript t is omitted, we use $N_\theta = \sum_{t=1}^T N_{t,\theta}$ to denote the total number of individuals of type θ. We additionally let $\bar{\rho}_{\ge t,\theta} = \frac{1}{T-t}\sum_{t'\ge t}\rho_{t,\theta}$ and similarly for $\bar{\sigma}^2_{\ge t,\theta}$ and $\bar{\mu}_{\ge t,\theta}$. A table with all our notation is provided in the online appendix.

3.1.3. Limitations and Extensions. The assumption that latent types Θ are finite is common in decision-making settings as, in practice, the set of possible types is approximated from historical data. One limiting assumption is that, in the online setting, the principal only knows the number of individuals from one location at a time. In reality, the principal could have some additional information about future locations, for example, via calling ahead, that could be incorporated in deciding an allocation. Our algorithmic approach naturally incorporates such additional information. Additionally, we assume a distinct set of individuals across each round and consider the rounds t as fixed and distinct locations.

3.2. Fairness and Efficiency in Off-line Allocations

To define an ex post fair allocation, that is, with a known number of individuals $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$ across rounds in $[T]$, we adopt an approach proposed by Varian (1974) (commonly referred to as "Varian fairness"), which is widely used in the operations research and economics literature. We refer to this as "fairness" for brevity; for a more detailed discussion on the advantages and limitations of this model, see Sugden (1984) or Online Appendix A.

Definition 1 (Fair Allocation). Given types Θ, a number of individuals of each type $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$, and utility functions $(u(\cdot,\theta))_{\theta\in\Theta}$, an allocation $X = \{X_{t,\theta} \in \mathbb{R}^K_+ | \sum_{t=1}^T \sum_{\theta\in\Theta} N_{t,\theta}X_{t,\theta} \le B\}$ is said to be fair if it simultaneously satisfies the following:

1. Envy-freeness (EF): For every pair of rounds t, t' and types θ, θ', we have $u(X_{t,\theta}, \theta) \ge u(X_{t',\theta'}, \theta)$.

2. Pareto efficiency (PE): For any allocation $Y \neq X$ such that $u(Y_{t,\theta}, \theta) > u(X_{t,\theta}, \theta)$ for some round t and type θ, there exists some other round t' and type θ' such that $u(Y_{t',\theta'}, \theta') < u(X_{t',\theta'}, \theta')$.

3. Proportional (Prop): For any round t, type θ, we have $u(X_{t,\theta}, \theta) \ge u(B/N, \theta)$, where $N = \sum_{t=1}^{T} \sum_{\theta\in\Theta} N_{t,\theta}$.

Whereas the three properties form natural desiderata for a fair allocation, the power of this definition lies in that asking for them to hold simultaneously rules out many natural (but unfair) allocation policies. In particular, allocation rules based on maximizing a global function, such as utilitarian welfare (sum of individual utilities) or egalitarian welfare (the maximin allocation or, more generally, the leximin allocation; Bogomolnaia and Moulin 2001, Lien et al. 2014, Manshadi et al. 2021) are Pareto efficient, but tend to violate individual envy-freeness as they focus on global optimality rather than per-individual guarantees. A remarkable exception to this, however, is the Nash social welfare, whose maximization leads to an allocation that is Pareto-efficient, envy-free, and proportional and, hence, fair.

Proposition 1 (Theorem 2.3 in Varian 1974). *For allocation X, the Nash social welfare is*

$$NSW(X) = \left(\prod_{t\in[T]} \prod_{\theta\in\Theta} u(X_{t,\theta}, \theta)^{N_{t,\theta}} \right)^{1/\sum_{t,\theta} N_{t,\theta}}. \tag{1}$$

Under linear utilities, an allocation X that maximizes NSW(X) is Pareto-efficient, envy-free, and proportional.

In addition to simultaneously ensuring PE, EF, and Prop properties, the NSW maximizing solution can also be efficiently computed via the following convex program called the EG program (Eisenberg 1961), obtained by taking the logarithm of the Nash social welfare:

$$\begin{aligned} \max_{X\in\mathbb{R}_+^{T\times\Theta\times K}} \quad & \sum_{t=1}^{T} \sum_{\theta\in\Theta} N_{t,\theta} \log(u(X_{t,\theta}, \theta)) \\ \text{s.t.} \quad & \sum_{t=1}^{T} \sum_{\theta\in\Theta} N_{t,\theta} X_{t,\theta} \le B. \end{aligned} \tag{2}$$

Important to note in our setting is that the optimal fair allocation in hindsight, which solves Equation (2) with a given number of individuals of each type across all rounds $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$, does not depend on round t. Indeed, any envy-free allocation can be formulated so $X_{t,\theta} = X_{t',\theta}$ (by setting $X_\theta = \frac{1}{t}\sum_t X_{t,\theta}$), and so we can instead consider the solution to

$$\max_{X\in\mathbb{R}_+^{\Theta\times K}} \sum_{\theta\in\Theta} N_\theta \log(u(X_\theta, \theta)) \qquad \text{s.t.} \sum_{\theta\in\Theta} N_\theta X_\theta \le B, \tag{3}$$

where we use $N_\theta = \sum_{t\in[T]} N_{t,\theta}$ to denote the total number of individuals across all rounds of type θ. The fact that the optimal solution in hindsight does not depend on the round t forms the basis for our algorithm Guarded-Hope.

3.3. Approximate Fairness and Efficiency in Online Allocations

Recall that, in our online setting, the principal allocates resources across each round in a fixed order $t = 1, \ldots, T$, whereupon, at round t, the principal sees $(N_{t,\theta})_{\theta\in\Theta}$ and decides on an allocation before continuing to the next round. A natural (albeit naive) approach in this setting could be to try and obtain allocations that satisfy Pareto-efficiency and envy-freeness on all sample paths. However, such an approach is not feasible even in the simplest online setting as the optimal solution in hindsight is often a unique function of the realized number of individuals across each round.

Proposition 2. *For $T = 2$ rounds, $|\Theta| = 1$ type, single resource, and linear utilities, for any nontrivial distribution $\mathcal{F}_2$, no online algorithm can guarantee ex post envy-freeness and Pareto-efficiency almost surely.*

Proof of Proposition 2. Let $\mathcal{F}_2 \sim 1 + \textsc{Bernoulli}(p)$ with $p \in (0, 1)$. For any value of N_1 with probability p, the optimal solution is $X^{opt} = B/(N_1 + 1)$, else $X^{opt} = B/(N_1 + 2)$. As any algorithm must decide how much to allocate at round $t = 1$ without knowledge of N_2, no algorithm can match the ex post fair solution almost surely. □

Proposition 2 shows that trying to simultaneously achieve ex post envy-freeness and Pareto-efficiency is futile, and hence, we need to consider approximate fairness notions. To this end, we define counterfactual envy, hindsight envy, and efficiency.

Definition 2 (Counterfactual Envy, Hindsight Envy, and Efficiency). Given individuals with types Θ, sizes $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$, and resource budgets $(B_k)_{k\in[K]}$, for any online allocation $(X_{t,\theta}^{alg})_{t\in[T],\theta\in\Theta} \in \mathbb{R}^k$, we define

- Counterfactual envy: The counterfactual distance of X^{alg} to envy-freeness as

$$\Delta_{EF} \triangleq \max_{t\in[T],\theta\in\Theta} \left\| u(X_{t,\theta}^{alg}, \theta) - u(X_{t,\theta}^{opt}, \theta) \right\|_\infty,$$

where X^{opt} is the optimal fair allocation in hindsight, that is, the solution to Equation (3) with true values $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$.

- Hindsight envy: The hindsight distance of X^{alg} to envy-freeness as

$$\textsc{Envy} \triangleq \max_{t,t'\in[T]^2,\theta,\theta'\in\Theta^2} u\left(X_{t',\theta'}^{alg}, \theta\right) - u\left(X_{t,\theta}^{alg}, \theta\right).$$

- Efficiency: The distance to efficiency as

$$\Delta_{efficiency} \triangleq \sum_{k\in K} \left(B_k - \sum_{t\in[T]} \sum_{\theta\in\Theta} N_{t,\theta} X_{t,\theta,k}^{alg} \right).$$

Our algorithm also provides ex post guarantees on hindsight proportionality defined via $\Delta_{prop} \triangleq \max_{t,\theta} u\left(\frac{B}{\sum_{t,\theta} N_{t,\theta}}, \theta\right) - u\left(X_{t,\theta}^{alg}, \theta\right)$.

These approximate fairness definitions are motivated by the problems faced by the FBST. Hindsight envy measures the algorithm's ability to ensure individuals are not envious of the allocations given to any other. Whereas this might serve as a natural first step toward a definition, an algorithm achieving low hindsight envy does not necessarily imply that individuals are eager to participate. In particular, the algorithm that allocates $X_{t,\theta} = 0$ for all t and θ trivially achieves hindsight envy of zero (suffering from large efficiency). Another consideration is ensuring allocations are close to what they should have been given based on observed information along the trajectory. Our measure of counterfactual envy addresses that, penalizing allocation algorithms based on how close they were at addressing individual's utility versus the optimal solution in hindsight. In fact, this metric is considered in the literature in a competitive ratio instead of additive sense (Friedman et al. 2015, 2017). Finally, efficiency is a natural yardstick for measuring an algorithm in order to ensure all of the resources that can be utilized are used.

We also note that all of these metrics are much stronger than the existing metrics in the literature because we provide hindsight guarantees that hold with high probability with respect to the distribution as opposed to weaker ex ante guarantees that only hold in expectation. Moreover, most approaches in the literature focus on defining a single optimization problem with a specified objective embodying "fairness" that attempts to capture desired goals. This is fundamentally flawed as the definition of the objective or choice of metric itself biases the outcomes toward a particular point along the trade-off curve between different criteria. The important challenge in this setting then is considering meaningful trade-offs between these metrics in the online setting (see Figure 1) and designing algorithms that achieve any point along the trade-off curve.

As highlighted earlier, the definition of counterfactual envy and efficiency are related. By using the fact that the optimal solution in hindsight X^{opt} is efficient, we can naively bound $\Delta_{efficiency}$ using Δ_{EF},

$$\Delta_{efficiency} = \sum_{t\in[T]}\sum_{\theta\in\Theta}\sum_{k\in K} N_{t,\theta}\left(X^{opt}_{t,\theta,k} - X^{alg}_{t,\theta,k}\right) \le \frac{TK\Delta_{EF}}{\|w\|_{min}}\sum_{\theta\in\Theta} N_\theta.$$

This naive bound is loose with unnecessary dependence on the number of locations T (see the counterfactual envy-efficiency uncertainty principle in Section 4).

4. Uncertainty Principles

In this section, we show parts (1), (2a), and (2b) from Theorem 1 concerning a lower bound on the achievable Δ_{EF} and the relationship between Δ_{EF}, Envy, and $\Delta_{efficiency}$ resulting from the envy-efficiency uncertainty principle. In all of these proofs, we consider the case of a single resource and single type, and assume that $u(X,\theta) = X$ for brevity and clarity in the presentation. However, the proofs extend directly to multiple resources in which one considers the setting with $|\Theta| = K$ and each type θ desires a unique resource.

We begin with part (1), the statistical uncertainty principle on the optimal fair allocation in hindsight, showing that no online algorithm is able to achieve counterfactual envy smaller than order $1/\sqrt{T}$. This arises because of the uncertainty in the number of individuals arriving in the future, forcing the algorithm to make a nontrivial decision on the allocation made to individuals in the first round.

Theorem 1 (Statistical Uncertainty Principle). *Let α be a constant with $\alpha + C\rho_{max}/\sigma^3_{min}\sqrt{T} < 1/2$, where C is an absolute constant. Then, with probability at least α, any online algorithm must incur*

$$\Delta_{EF} \ge \beta_{avg}\frac{3\Phi^{-1}\left(1-\alpha-\frac{C\rho_{max}}{\sigma^3_{min}\sqrt{T}}\right)\sigma_{min}}{4\sqrt{T}}.$$

Proof of Theorem 1. We use the generalized Berry–Esseen theorem (Berry 1941). Recall that, for all t, $\mathrm{Var}[N_t] = \sigma_t > 0$ and $\mathbb{E}[|N_t - \mathbb{E}[N_t]|] = \rho_t < \infty$, and moreover, $X^{opt}_t = B/N$ for all t, where $N = \sum_{t\in[T]} N_t$. Let us denote $\bar\sigma^2 = \frac{1}{T}\sum_{t\in[T]}\sigma_t^2$ and $\bar\rho = \frac{1}{T}\sum_{t\in[T]}\rho_t$ and let Φ be the cumulative distribution function of a standard normal. Using Berry–Esseen, it holds that, for an absolute constant C, for all $z\in\mathbb{R}$,

$$\Phi(z) - \frac{C\bar\rho}{\bar\sigma^3\sqrt{T}} \le \mathbb{P}\left(X_{opt} \ge \frac{B}{\mathbb{E}[N] + z\bar\sigma\sqrt{T}}\right) \le \Phi(z) + \frac{C\bar\rho}{\bar\sigma^3\sqrt{T}}.$$

Taking $z = -y$ and using the lower bound, we have that, with probability at least $\Phi(-y) - \frac{C\bar\rho}{\bar\sigma^3\sqrt{T}}$,

$$X^{opt} \ge \frac{B}{\mathbb{E}[N] - y\bar\sigma\sqrt{T}} \ge \frac{B}{\mathbb{E}[N]}\left(1 + \frac{y\bar\sigma\sqrt{T}}{\mathbb{E}[N]}\right).$$

Taking $z = y$ and using the upper bound, we have that, with probability at least $\Phi(-y) - \frac{C\bar\rho}{\bar\sigma^3\sqrt{T}}$,

$$X^{opt} \le \frac{B}{\mathbb{E}[N] + y\bar\sigma\sqrt{T}} \le \frac{B}{\mathbb{E}[N]}\left(1 - \frac{y\bar\sigma\sqrt{T}}{2\mathbb{E}[N]}\right).$$

Note that these intervals are nonoverlapping for $y > 0$. As the algorithm must decide on a value X^{alg}_1 to allocate for the first round, then with probability at least $\Phi(-y) - \frac{C\bar\rho}{\bar\sigma^3\sqrt{T}} \ge \Phi(-y) - \frac{C\rho_{max}}{\sigma^3_{min}\sqrt{T}}$,

$$\|X^{alg} - X^{opt}\|_\infty \ge \min_{x\in\mathbb{R}}\max\left(\left|\frac{B}{\mathbb{E}[N]}\left(1 - \frac{y\bar\sigma\sqrt{T}}{2\mathbb{E}[N]}\right) - x\right|, \left|\frac{B}{\mathbb{E}[N]}\left(1 + \frac{y\bar\sigma\sqrt{T}}{\mathbb{E}[N]}\right) - x\right|\right) = \frac{B}{\mathbb{E}[N]}\frac{3y\bar\sigma\sqrt{T}}{4\mathbb{E}[N]}.$$

Taking $y = \Phi^{-1}\left(1 - \alpha - \frac{C\rho_{max}}{\sigma_{min}^3\sqrt{T}}\right)$, which is positive, then we get with probability at least α that

$$\|X^{alg} - X^{opt}\|_\infty \geq \frac{B}{\mathbb{E}[N]} \frac{3\Phi^{-1}\left(1 - \alpha - \frac{C\rho_{max}}{\sigma_{min}^3\sqrt{T}}\right)\bar{\sigma}}{4\mathbb{E}[N]}$$
$$\geq \beta_{avg} \frac{3\Phi^{-1}\left(1 - \alpha - \frac{C\rho_{max}}{\sigma_{min}^3\sqrt{T}}\right)\sigma_{min}}{4\sqrt{T}}. \quad \square$$

We next show the first part of the envy-efficiency uncertainty principle (2a), highlighting that any online algorithm that achieves a factor of L_T on counterfactual envy necessarily suffers efficiency of at least $1/L_T$. This result follows from the statistical uncertainty in the number of individuals arriving in the final L_T^{-2} rounds and the fact that ensuring a bounded envy requires any online algorithm to save enough budget to allocate a minimum allocation to all future arriving individuals.

Theorem 2 (Counterfactual Envy-Efficiency Uncertainty Principle). *Let $\alpha < 1/8$ be a constant such that $3\alpha + C\rho_{max}/\sigma_{min}^3\sqrt{T} < 1/2$ for an absolute constant C. Any online algorithm that achieves $\Delta_{EF} \leq L_T = o(1)$ with probability at least $1 - \alpha$ must also incur waste $\Delta_{efficiency} \geq \tilde{C}\ \min\{\sqrt{T}, 1/L_T\}$, where*

$$\tilde{C} = (\beta_{avg} - o(1))^2 \frac{\Phi^{-1}\left(1 - 3\alpha - \frac{C\rho_{max}}{\sigma_{min}^3\sqrt{T}}\right)^2 \sigma_{min}^2}{24\sqrt{2\rho_{max}^2 \log(T/\alpha)}\mu_{max}}$$

with probability at least $1/12 - o(1)$.

Proof of Theorem 2. In order for the algorithm to guarantee that $\Delta_{EF} = \|X^{alg} - X^{opt}\|_\infty \leq L_T$ with probability at least $1 - \alpha$, it must limit all allocations made to the interval $[\frac{B}{N} - L_T, \frac{B}{N} + L_t]$ as $X^{opt} = B/N$. Moreover, a straight forward application of Hoeffding's inequality shows that $|N - \mathbb{E}[N]| \leq \tilde{c}\sqrt{T}$ with probability $1 - \alpha$, where $\tilde{c} = \sqrt{2\rho_{max}^2 \log(T/\alpha)}$. Using this, algebraic manipulations, and simplifying, one can show that the following event:

$$\mathcal{D} = \bigcap_{t \in [T]} \left\{ X_t^{alg} \in \left[\frac{B}{\mathbb{E}[N]} - \frac{\tilde{c}\sqrt{T}}{\mathbb{E}[N]} - L_T, \frac{B}{\mathbb{E}[N]} + \frac{2\tilde{c}\sqrt{T}}{\mathbb{E}[N]} + L_T\right]\right\}$$

occurs with probability at least $1 - 2\alpha$. We interpret these lower and upper thresholds on allocations made by the algorithm as guardrails.

Recall that we use the notation B_t^{alg} to denote the budget remaining for the algorithm at the start of round t. We begin by defining three events for a fixed round $t \leq T$ and constant $z = \Phi^{-1}\left(1 - 3\alpha - \frac{C\rho_{max}}{\sigma_{min}^3\sqrt{T}}\right) > 0$:

$$\mathcal{A} = \{N_{\geq t} \leq \mathbb{E}[N_{\geq t}]\}$$
$$\mathcal{B} = \{N_{\geq t} \geq \mathbb{E}[N_{\geq t}] + z\bar{\sigma}_{\geq t}\sqrt{T - t + 1}\}$$
$$\mathcal{C} = \left\{ B_t^{alg} \geq \left(\frac{B}{\mathbb{E}[N]} - \frac{\tilde{c}\sqrt{T}}{\mathbb{E}[N]} - L_T\right) (\mathbb{E}[N_{\geq t}] + z\bar{\sigma}_{\geq t}\sqrt{T - t + t})\right\}.$$

By the Berry–Esseen theorem (Berry 1941), we know that $\mathbb{P}(\mathcal{A}) \geq \frac{1}{2} - \frac{C\bar{\rho}_{\geq t}}{\bar{\sigma}_{\geq t}^3\sqrt{T-t+1}} \geq \frac{1}{2} - \frac{C\rho_{max}}{\sigma_{min}^3\sqrt{T}}$ and $\mathbb{P}(\mathcal{B}) \geq \Phi(-z) - \frac{C\bar{\rho}_{\geq t}}{\bar{\sigma}_{\geq t}^3\sqrt{T-t+1}} \geq 3\alpha$ by choice of z.

We first show that $\neg\mathcal{C} \cap \mathcal{B}$ implies $\neg\mathcal{D}$ (or, equivalently, by taking the contrapositive that $\mathcal{D}$ implies that $\mathcal{B}$ implies $\mathcal{C}$), which gives us that $\mathbb{P}(\neg\mathcal{C} \cap \mathcal{B}) \leq \mathbb{P}(\neg\mathcal{D})$. These two conditions ($\mathcal{D}$ and $\mathcal{B}$) dictate that the algorithm must have a lot of budget by allocating within the guardrails based on the number of individuals arriving in the future being small. Indeed, under events $\mathcal{D}$ and $\mathcal{B}$ we have

$$B_t^{alg} \geq \sum_{t' \geq t} N_{t'} X_{t'}^{alg} \geq N_{\geq t}\left(\frac{B}{\mathbb{E}[N]} - \frac{\tilde{c}\sqrt{T}}{\mathbb{E}[N]} - L_T\right)$$
$$\geq (\mathbb{E}[N_{\geq t}] + z\bar{\sigma}_{\geq t}\sqrt{T - t + 1})\left(\frac{B}{\mathbb{E}[N]} - \frac{\tilde{c}\sqrt{T}}{\mathbb{E}[N]} - L_T\right).$$

Moreover, because the algorithm is nonanticipatory, we know that the events $\mathcal{C}$ and $\mathcal{B}$ are independent. Thus, we have that $\mathbb{P}(\neg\mathcal{C} \cup \mathcal{B}) = \mathbb{P}(\neg\mathcal{C})\mathbb{P}(\mathcal{B})$. Using the bound on $\mathbb{P}(\mathcal{B})$ and the fact that $\mathbb{P}(\neg\mathcal{C} \cup \mathcal{B}) \leq \mathbb{P}(\neg\mathcal{D}) \leq 2\alpha$, we get that $\mathbb{P}(\neg\mathcal{C}) \leq \frac{2}{3}$.

Now, we consider the event $\mathcal{C} \cap \mathcal{A} \cap \mathcal{D}$. Using that the allocations must be bounded by event $\mathcal{D}$, we have that the waste is at least

$$\Delta_{efficiency} = B - \sum_{i=1}^{T} X_i^{alg} N_i = B_t^{alg} - \sum_{i \geq t} X_i^{alg} N_i$$
$$\geq (\mathbb{E}[N_{\geq t}] + z\bar{\sigma}_{\geq t}\sqrt{T - t + 1})\left(\frac{B}{\mathbb{E}[N]} - \frac{\tilde{c}\sqrt{T}}{\mathbb{E}[N]} - L_T\right)$$
$$- \left(\frac{B}{\mathbb{E}[N]} + \frac{2\tilde{c}\sqrt{T}}{\mathbb{E}[N]} + L_T\right)\mathbb{E}[N_{\geq t}]$$
$$\geq \left(\frac{B}{\mathbb{E}[N]} - \frac{\tilde{c}\sqrt{T}}{\mathbb{E}[N]} - L_T\right) z\bar{\sigma}_{\geq t}\sqrt{T - t + 1}$$
$$- 3\left(L_T + \frac{\tilde{c}\sqrt{T}}{\mathbb{E}[N]}\right)\mathbb{E}[N_{\geq t}].$$

The inequality follows from lower bounding B_t^{alg} with the amount required to be reserved up to location t (i.e., event $\mathcal{C}$), and upper bounding the maximum amount of budget that can be expended for locations $i \geq t$ when $N_{\geq t} \leq \mathbb{E}[N_{\geq t}]$ (i.e., event $\mathcal{A}$).

Recall that $\mathbb{E}[N_{\geq t}] = (T - t + 1)\bar{\mu}_{\geq t}$ so that, whereas the first term increases with $(T - t + 1)$, the second term decreases with $(T - t + 1)$. Solving for the maximum value in terms of t yields

$$\Delta_{efficiency} \geq \frac{1}{12}\left(\frac{B}{\mathbb{E}[N]} - \frac{\tilde{c}\sqrt{T}}{\mathbb{E}[N]} - L_T\right)^2 \frac{z^2\bar{\sigma}_{\geq t}^2}{\left(L_T + \frac{\tilde{c}\sqrt{T}}{\mathbb{E}[N]}\right)\bar{\mu}_{\geq t}}$$
$$\geq \frac{1}{12}\left(\frac{B}{\mathbb{E}[N]} - \frac{\tilde{c}\sqrt{T}}{\mathbb{E}[N]} - L_T\right)^2 \frac{z^2\sigma_{min}^2}{(L_T + \tilde{c}/\sqrt{T})\mu_{max}}.$$

The probability of this event is lower bounded by $\mathbb{P}(\mathcal{C} \cap \mathcal{A} \cap \mathcal{D}) \geq \mathbb{P}(\mathcal{C} \cap \mathcal{A}) - \mathbb{P}(\neg\mathcal{D}) \geq \mathbb{P}(\mathcal{C})\mathbb{P}(\mathcal{A}) - 2\alpha \geq (1 - \frac{2}{3})\left(\frac{1}{2} - \frac{C\rho_{max}}{\sigma_{min}^3\sqrt{T}}\right) - 2\alpha \geq \frac{1}{12} - o(1)$. Plugging in the value of z and simplifying terms yields the final result. □

Finally we show the second part of the envy-efficiency uncertainty principle (2b), highlighting that any online algorithm that achieves a factor of L_T on hindsight envy necessarily suffers efficiency of at least $\min\{\sqrt{T}, 1/L_T\}$. This result follows from the previous lower bound (2a), combined with an almost sure relationship between Δ_{EF} and ENVY. We start with this brief lemma relating the two notions of envy.

Lemma 1 (Relation Between Envy and Δ_{EF}). *For any valid online allocation algorithm, we have the following almost surely:*

$$\Delta_{EF} - \frac{1}{N}\Delta_{efficiency} \leq \text{ENVY} \leq 2\Delta_{EF}.$$

Proof of Lemma 1. The upper bound follows immediately from applying the triangle inequality around $X^{opt} = \frac{B}{N}$. For the lower bound, we instead show that $\frac{1}{N}\Delta_{efficiency} \geq \Delta_{EF} - \text{ENVY}$. Here, we set $L = \min_i X_i$ and $U = \max_i X_i$ to be the maximum and minimum allocations given out by the algorithm. Note that $\text{ENVY} = U - L$ and $\Delta_{EF} = \max\left\{\left|\frac{B}{N} - U\right|, \left|\frac{B}{N} - L\right|\right\}$. First notice no algorithm can have $L > \frac{B}{N}$ because of the feasibility of allocations made by the algorithm. Thus, we get that $\Delta_{EF} = \max\left\{\left|\frac{B}{N} - U\right|, \frac{B}{N} - L\right\}$. We show the inequality breaking into cases on the side that achieves the max.

Case 1: $\Delta_{EF} = \frac{B}{N} - L$.
In this setting, we have that $\Delta_{EF} = \frac{B}{N} - L$, $\text{ENVY} = U - L$.
Using this, we can show

$$\begin{aligned}\Delta_{efficiency} &= B - \sum_i N_i X_i = \sum_i N_i\left(\frac{B}{N} - X_i\right) \\ &\geq \sum_i N_i\left(\frac{B}{N} - U\right) = N\left(\frac{B}{N} - U\right) \\ &= N(\Delta_{EF} - \text{ENVY}).\end{aligned}$$

Case 2: $\Delta_{EF} = |\frac{B}{N} - U|$.
This implies that $L \leq \frac{B}{N} \leq U$ as otherwise the maximum would be achieved by $\frac{B}{N} - L$. Thus, we get that $\Delta_{EF} - \text{ENVY} = U - \frac{B}{N} - (U - L) = L - \frac{B}{N}$, which is negative, so the inequality is trivially true. □

Using this, we are able to show (2b) in the envy-efficiency uncertainty principle, relating the necessary trade-off between ENVY and $\Delta_{efficiency}$.

Theorem 3 (Hindsight Envy-Efficiency Uncertainty Principle). *Let $\alpha < \frac{1}{8}$ be a constant such that $3\alpha + \frac{C\rho_{max}}{\sigma_{min}^3\sqrt{T}} < \frac{1}{2}$ for an absolute constant C. Any online algorithm that achieves $\text{ENVY} \leq L_T = o(1)$ with probability at least $1 - \alpha$ must also incur waste $\Delta_{efficiency} \geq \tilde{C}\min\{\sqrt{T}, 1/L_T\} - o(1)$, where $\tilde{C}$ is as in Theorem 2 with probability at least $\frac{1}{12} - o(1)$.*

Proof of Theorem 3. Suppose the online algorithm achieves $\text{ENVY} \leq L_T$ with probability at least $1 - \alpha$. However, using Lemma 1, we get that $\Delta_{EF} - \frac{1}{N}\Delta_{efficiency} \leq \text{ENVY} \leq L_T$. Hence, we have that $\Delta_{EF} \leq L_T + \frac{1}{N}\Delta_{efficiency}$ with probability at least $1 - \alpha$. Denote by $\tilde{C}$ as the terms on the right-hand side of Theorem 2, and applying the result there for the case when the $1/L_T$ term attains the minimum, we get that $\Delta_{efficiency} \geq \tilde{C}\frac{1}{L_T + \frac{1}{N}\Delta_{efficiency}}$. Rearranging the inequality gives us that $\Delta_{efficiency} \geq \frac{N}{2}\left[\sqrt{\frac{4\tilde{C}}{N} + L_T^2} - L_T\right]$. The final bound comes from taking the first term of the Taylor series about infinity with the additional $o(1)$ factor. □

5. Sensitivity and Concentration on Counterfactual Optimal Fair Allocation

The lower bounds presented in Section 4 highlight a key facet of algorithm design in this setting: generating lower and upper guardrail allocations. Suppose we were able to construct envy-free allocations $(\underline{X}_\theta)_{\theta\in\Theta}$ and $(\overline{X}_\theta)_{\theta\in\Theta}$ such that $\|\overline{X}_\theta - \underline{X}_\theta\|_\infty \lesssim L_T$ for a given parameter L_T. If the algorithm was able to ensure that all of the allocations made to individuals of type θ are within $[\underline{X}_\theta, \overline{X}_\theta]$, it is not difficult to show that $\text{ENVY} \lesssim L_T$. However, if we additionally desire a bound of $\Delta_{EF} \lesssim L_T$, the same philosophy requires that we are able to establish that, with high probability,

$$u(\underline{X}_\theta, \theta) \leq u(X_\theta^{opt}, \theta) \leq u(\overline{X}_\theta, \theta) \qquad \forall\theta \in \Theta \text{ with } L_T \gtrsim 1/\sqrt{T}. \tag{4}$$

Motivated by these two use cases, we turn our attention to sensitivity and concentration properties on solutions to the Eisenberg–Gale program. Unfortunately, the true EG program for the counterfactual optimal fair allocation depends on the unknown vector of number of individuals of each type $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$. As such, our algorithms are motivated by solving information-relaxed versions of the EG program, appealing to sensitivity and concentration on the optimizers of the program instead of the objective value as is typically done in competitive ratio analysis.

For the time being, we assume that we are given concentration inequalities of the following form: with probability at least $1 - \delta$, we have that, for every t and θ, $|\mathbb{E}[N_{>t,\theta}] - N_{>t,\theta}| \leq \text{CONF}_{t,\theta}$. As this concentration only depends on the assumptions on the variables $N_{t,\theta}$, we include a simple form of $\text{CONF}_{t,\theta}$ scaling as $\sqrt{T - t}$ using Hoeffding's inequality in Lemma EC.3, but see Section 7 for extensions.

Consider the Eisenberg–Gale program from Section 3 with multiple types θ and K resources as specified in Equation (3). Recall that the dual variables corresponding to the budget feasibility constraint p_k can be thought of as prices for the corresponding resources (Nisan et al. 2007). We start with a lemma showing properties of the optimal solution to the Eisenberg–Gale program with various numbers of individuals of each type vector $(N_\theta)_{\theta\in\Theta}$.

Lemma 2 (Sensitivity of Solutions to the Eisenberg–Gale Program). *Let $x((N_\theta)_{\theta\in\Theta})$ and $p((N_\theta)_{\theta\in\Theta})$ denote the optimal primal and dual solution to the Eisenberg–Gale program (Equation (3)) for a given vector of individuals of each type $(N_\theta)_{\theta\in\Theta}$. Then, we have*

1. *Scaling: If $\tilde{N}_\theta = (1+\zeta)N_\theta$ for every $\theta \in \Theta$ and $\zeta \geq 0$, then we have that*

$$x((\tilde{N}_\theta)_{\theta\in\Theta}) = \frac{x((N_\theta)_{\theta\in\Theta})}{1+\zeta}$$

$$p((\tilde{N}_\theta)_{\theta\in\Theta}) = (1+\zeta)p((N_\theta)_{\theta\in\Theta})$$

$$u(x((N_\theta)_{\theta\in\Theta})_\theta, \theta) - u(x((\tilde{N}_\theta)_{\theta\in\Theta})_\theta, \theta) = \left(1 - \frac{1}{1+\zeta}\right) \max_k \frac{w_{\theta,k}}{p((N_\theta)_{\theta\in\Theta})_k}.$$

2. *Monotonicity: If $N_\theta \leq \tilde{N}_\theta$ for every $\theta \in \Theta$, then we have*

$$p((\tilde{N}_\theta)_{\theta\in\Theta}) \geq p((N_\theta)_{\theta\in\Theta})$$
$$u(x((\tilde{N}_\theta)_{\theta\in\Theta})_\theta, \theta) \leq u(x((N_\theta)_{\theta\in\Theta})_\theta, \theta) \quad \forall \theta \in \Theta.$$

These Lipschitz properties follow from the Fisher market interpretation of the Eisenberg–Gale optimum, which corresponds to market-clearing allocations in a setting with $|\Theta|$ agents, each with an endowment or budget of N_θ. The second property is a generalization of the competitive monotonicity property Devanur et al. (2002). See Online Appendix E for the full proof.

Recall that our goal is to construct lower and upper threshold allocations about $X^{opt} = x((N_\theta)_{\theta\in\Theta})$, where $N_\theta = \sum_t N_{t,\theta}$ is the true (random) number of individuals of type θ arriving over all rounds. First suppose we were able to construct $\overline{n}_\theta \geq N_\theta$ for all θ and set $\underline{n}_\theta = (1-\gamma)\overline{n}_\theta$ for some constant γ chosen based on L_T. Setting the guardrails as $\underline{X}_\theta = x((\overline{n}_\theta)_{\theta\in\Theta})$ and $\overline{X}_\theta = x((\underline{n}_\theta)_{\theta\in\Theta})$ are envy free (by Proposition 1) and, for the correct value of γ, satisfy the bounds needed to ensure $\text{Envy} \lesssim L_T$ (by part (1) of Lemma 2). However, for a large enough γ (or, equivalently, a large enough L_T), we can additionally ensure that $\underline{n}_\theta \leq N_\theta$ and appeal to the monotonicity property (2) of Lemma 2 to ensure Equation (4). This assumption includes not only i.i.d. demands, but also demand distributions that arise from Markov chains, latent variable models, or known cumulative sums from different Chernoff style arguments (see Section 7 for more details). The following lemma shows the final construction of our guardrails by appropriately choosing $\overline{n}_\theta$ and $\underline{n}_\theta$ and appealing to Lemma 2.

Theorem 4 (Construction of Guardrail Allocations). *Let $X^{opt} = x((N_\theta)_{\theta\in\Theta})$ denote the optimal solution to the Eisenberg–Gale program for a given vector of individuals of each type $(N_\theta)_{\theta\in\Theta}$. Further suppose that, with probability at least $1-\delta$, we have for all $\theta \in \Theta$ $|N_\theta - \mathbb{E}[N_\theta]| \leq \text{Conf}_\theta$. Given any $L_T \geq 0$ and setting*

$$\overline{n}_\theta = \mathbb{E}[N_\theta]\left(1 + \max_\theta \frac{\text{Conf}_\theta}{\mathbb{E}[N_\theta]}\right)$$
$$\underline{n}_\theta = \mathbb{E}[N_\theta](1-c) \quad \text{for}$$
$$c = \frac{\|w\|_{min}\|\beta_{avg}\|_{min}}{\|w\|_\infty^2} L_T \left(1 + \max_\theta \frac{\text{Conf}_\theta}{\mathbb{E}[N_\theta]}\right) - \max_\theta \frac{\text{Conf}_\theta}{\mathbb{E}[N_\theta]},$$

then almost surely we have that

1. $u(x((\underline{n}_\theta)_{\theta\in\Theta})_\theta, \theta) - u(x((\overline{n}_\theta)_{\theta\in\Theta})_\theta, \theta) \leq L_T$.
2. $\|x((\overline{n}_\theta)_{\theta\in\Theta}) - x((\underline{n}_\theta)_{\theta\in\Theta})\|_\infty \geq L_T \frac{\|\beta_{avg}\|_{min}^2 \|w\|_{min}}{\|w\|_\infty}$.
3. $\|x((\overline{n}_\theta)_{\theta\in\Theta}) - x((\underline{n}_\theta)_{\theta\in\Theta})\|_\infty \leq L_T \frac{\|B\|_\infty \|\beta_{avg}\|_{min}\|w\|_{min}}{\|w\|_\infty}$.

If, in addition, $L_T \geq 2\frac{\|w\|_\infty^2}{\|w\|_{min}\|\beta_{avg}\|_{min}} \max_\theta \frac{\text{Conf}_\theta}{\mathbb{E}[N_\theta]}$, then with probability at least $1-\delta$, we have

4. $\underline{n}_\theta \leq N_\theta \leq \overline{n}_\theta$,
5. $u(x((\overline{n}_\theta)_{\theta\in\Theta})_\theta, \theta) \leq u(X_\theta^{opt}, \theta) \leq u(x((\overline{n}_\theta)_{\theta\in\Theta})_\theta, \theta)$.

See Online Appendix E for the full proof. Using a straightforward application of Hoeffding's inequality, we notice that this construction ensures that we are able to guarantee a bound of L_T on the difference in utilities for any $L_T \geq 2\frac{\|w\|_\infty^2}{\|w\|_{min}\|\beta_{avg}\|_{min}} \sqrt{\frac{2\rho_{max}^2 \log(T|\Theta|/\delta)}{T}}$.

6. Guarded-Hope

Here, we define our algorithm Guarded-Hope. The algorithm takes as input a budget B, expected number of each type $(\mathbb{E}[N_\theta])_{\theta\in\Theta}$, confidence terms $(\text{Conf}_{t,\theta})_{\theta\in\Theta}$, and a desired bound L_T on the Δ_{EF} and Envy. Assuming the lower and upper threshold allocations are constructed such that we can guarantee the results from Section 5, our algorithm is able to achieve any envy-efficiency trade-off as developed in Theorem 1. The algorithm relies on two main components, both of which we believe to be necessary in developing an algorithm to achieve the envy-efficiency uncertainty principle (as removing any one of them leads to breakdowns as is discussed in Section 7). We start by describing the high-level ideas needed in the algorithm before describing the pseudocode (with full algorithm description in Algorithm 1). The proof that Guarded-Hope achieves the desired bounds is deferred to Section 7.

6.1. Guardrails on Optimal Fair Allocation in Hindsight

As a result of Theorem 1, we see that no online algorithm can guarantee $\Delta_{EF} \lesssim \frac{1}{\sqrt{T}}$. Moreover, the proof highlights that any algorithm that satisfies a bound on Δ_{EF} or $\textsc{Envy} \lesssim L_T$ must limit allocations based on guardrails with high probability. As such, our algorithm uses the construction from Section 5 to obtain estimates $\overline{X} = x((\underline{n}_\theta)_{\theta\in\Theta})$ and $\underline{X} = x((\overline{n}_\theta)_{\theta\in\Theta})$, which are both envy-free and satisfy that $\max_{t,\theta}|u(\overline{X}_{t,\theta},\theta) - u(\underline{X}_{t,\theta},\theta)| \le L_T$. If, in addition, $L_T \gtrsim 1/\sqrt{T}$, we have that $u(\underline{X}_\theta,\theta) \le u(X_\theta^{opt},\theta) \le u(\overline{X}_\theta,\theta)$.

The allocations $\overline{X}_\theta$ and $\underline{X}_\theta$ are used by the algorithm as guardrails, for which all allocations made by the algorithm for a type θ are forced to fall within $\{\underline{X}_\theta, \overline{X}_\theta\}$. With this requirement, on sample paths when we do not run out of budget, then we trivially have an upper bound on $\textsc{Envy} \lesssim L_T$ and $\Delta_{EF} \le \max\{1/\sqrt{T}, L_T\}$. Thus "accepting" the first round loss in envy-freeness allows us to limit all future allocations to the guardrails generated by that uncertainty.

6.2. Minimizing Waste via Online Stochastic Packing

Once the guardrails $\overline{X}_\theta$ and $\underline{X}_\theta$ are found to verify a bound on the approximate envy up to a factor of L_T, we change the focus to instead try and minimize the loss of efficiency. Thanks to our guardrails, we develop the algorithm to match a clairvoyant benchmark policy that minimizes the resource waste with the knowledge of $(N_{t,\theta})_{t\in[T],\theta\in\Theta}$ and simultaneously limits the allocations to lie between $\{\underline{X}_\theta, \overline{X}_\theta\}$. This can be thought of as solving an online stochastic packing problem (whose objective is to minimize efficiency loss) with the addition of guardrail constraints (i.e., our minimum and maximum allocation constraints). In this setting, after writing down the stochastic packing optimization program, a competitive algorithm arises naturally by ensuring that the budget remaining for the algorithm is enough to satisfy a high-probability bound on the resources required to allocate $\underline{X}$ to every individual arriving in the future. This idea is formalized in Section 7 and takes motivation in recent developments on Bayesian prophet benchmarks for online bin packing problems (Vera and Banerjee 2019).

6.3. Algorithm Description

Let $B_{t,k}^{alg}$ denote the budget remaining to the principal for resource k at iteration t, that is, $B_k - \sum_{t'<t}\sum_\theta N_{t',\theta}X_{t',\theta,k}$. Assume the algorithm is given the expected demands $(\mathbb{E}[N_\theta])_{\theta\in\Theta}$ and confidence terms $\textsc{Conf}_{t,\theta}$ such that $|N_{>t,\theta} - \mathbb{E}[N_{>t,\theta}]| \le \textsc{Conf}_{t,\theta}$ with high probability. Let $\gamma = \max_\theta \frac{Conf_{0,\theta}}{\mathbb{E}[N_{\ge 1,\theta}]}$. Given a desired bound on envy L_T, the algorithm computes the guardrails by

$$\underline{X} = x((\overline{n}_\theta)_{\theta\in\Theta}) \text{ for } \overline{n}_\theta = \left(1 + \max_\theta \frac{\textsc{Conf}_\theta}{\mathbb{E}[N_\theta]}\right)\mathbb{E}[N_\theta]$$
$$\overline{X} = x((\underline{n}_\theta)_{\theta\in\Theta}) \text{ for } \underline{n}_\theta = (1-c)\mathbb{E}[N_\theta]$$
$$\text{for } c = \frac{\|w\|_{min}\|\beta_{\text{avg}}\|_{min}}{\|w\|_\infty^2} \mathrm{L_T}\left(1 + \max_\theta \frac{\textsc{Conf}_\theta}{\mathbb{E}[N_\theta]}\right) - \max_\theta \frac{\textsc{Conf}_\theta}{\mathbb{E}[N_\theta]}.$$

Here, $x(\cdot)$ denotes the solution to Equation (3). Note that, as long as $L_T \gtrsim \sqrt{\log(|\Theta|T/\delta)/T}$, the utility of the optimal allocation is sandwiched by the utilities of $\underline{X}$ and $\overline{X}$ according to Section 5.

Our algorithm allocates to type θ according to these thresholds $\underline{X}_\theta$ and $\overline{X}_\theta$ in order to ensure the guarantee of Δ_{EF} of at most L_T, and simultaneously trying to eliminate as much waste as possible. At each time t, for each resource $k \in [K]$,

1. (Insufficient budget) If $B_{t,k}^{alg} \le N_{t,\theta}\underline{X}_{\theta,k}$, then divide the resources equally among all remaining individuals for this round.
2. (Sufficient budget to promise lower threshold) If $B_{t,k}^{alg} \ge \overline{X}_{\theta,k}N_{t,\theta} + \underline{X}_{\theta,k}(\mathbb{E}[N_{>t,\theta}] + \textsc{Conf}_{t,\theta})$, then set $X_{t,\theta,k}^{alg} = \overline{X}_{\theta,k}$ for each $\theta \in \Theta$.
3. (Insufficient budget to promise lower threshold) Otherwise, set $X_{t,\theta,k}^{alg} = \underline{X}_{\theta,k}$ for each $\theta \in \Theta$.

Our algorithm is easy to implement in practice; in particular, it requires solving the Eisenberg–Gale program Equation (3) only twice to obtain $\underline{X}$ and $\overline{X}$. In contrast, other versions of certainty equivalence algorithms require frequent resolves of the Eisenberg–Gale (see Online Appendix B). This allows Guarded-Hope to scale easily to multiple resources and a larger number of types (as it only involves solving for the "optimistic" and "pessimistic" allocation rules, which are done off-line with historical data). Moreover, it allows practitioners to leverage work on poly-time algorithms for solving the Eisenberg–Gale program (Devanur et al. 2002). It also extends easily to more complex information structures (see Section 7 for a discussion).

7. Envy and Efficiency Bound for Guarded-Hope

We are now ready to show the bound on Δ_{EF}, Envy and $\Delta_{efficiency}$ for Guarded-Hope, relying on the construction of $\overline{X}_\theta$ and $\underline{X}_\theta$ from Section 5. We note that these guarantees match the envy-efficiency uncertainty principles from Section 4 up to problem-dependent constants and logarithmic terms in T. Afterward, we comment on extending the distributional assumptions on $N_{t,\theta}$ to more robust settings.

Theorem 5. *Given budget B, expected number of types $(\mathbb{E}[N_\theta])_{\theta\in\Theta}$, and confidence terms $(\textsc{Conf}_{t,\theta})_{\theta\in\Theta}$ such that, with probability at least $1-\delta$, $|N_{>t,\theta} - \mathbb{E}[N_{>t,\theta}]| \le \textsc{Conf}_{t,\theta}$ for all $t \in [T]$, $\theta \in \Theta$, Guarded-Hope with parameter L_T is able to achieve with probability at least $1-\delta$ (where $\lesssim$ drops poly-logarithmic factors of T, $o(1)$ terms, and absolute constants)*

$$\textsc{Envy} \le L_T \qquad \Delta_{EF} \lesssim \max\{1/\sqrt{T}, L_T\}$$
$$\Delta_{efficiency} \lesssim \min\{\sqrt{T}, 1/L_T\} \qquad \Delta_{prop} \lesssim \max\{1/\sqrt{T}, L_T\}.$$

Proof of Theorem 5. Define the event $\mathcal{E} = \{\forall t \in [T], \forall \theta \in \Theta : |N_{\geq t,\theta} - \mathbb{E}[N_{\geq t,\theta}]| \leq \textsc{Conf}_{t-1,\theta}\}$. By assumption, we know that $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$. The following lemma shows that the algorithm ensures it has enough budget to allocate according to the lower threshold for everyone arriving in the future. Recall that $B_{t,k}^{alg}$ denotes the budget remaining at the start of round t for resource k.

Lemma 3. *Under the event $\mathcal{E}$, for every resource k and time $t \in [T]$ it follows that*

$$B_{t,k}^{alg} \geq \sum_{\theta \in \Theta} N_{\geq t,\theta} \underline{X}_{\theta,k}.$$

As a result, the algorithm is able to guarantee that, at every iteration, $X_{t,\theta,k}^{alg}$ is either $\overline{X}_{\theta,k}$ or $\underline{X}_{\theta,k}$.

Proof of Lemma 3. We show the first statement by induction on t. The second statement follows immediately.

Base Case $t = 1$
Here, we have that $B_1^{alg} = B$, and by construction of $\underline{X}_\theta$, we have that

$$B_{1,k}^{alg} \geq \sum_{\theta \in \Theta} \overline{n}_\theta \underline{X}_{\theta,k} \text{ (by feasibility) } \geq \sum_{\theta \in \Theta} N_\theta \underline{X}_{\theta,k} \text{ (by event } \mathcal{E}).$$

Step Case $t-1 \to t$
We split into two cases based on the allocation. If $X_{t-1,\theta,k}^{alg} = \underline{X}_{\theta,k}$, then by the induction hypothesis, $B_{t,k}^{alg} = B_{t-1,k}^{alg} - \sum_{\theta \in \Theta} N_{t-1,\theta} \underline{X}_{\theta,k} \geq \sum_{\theta \in \Theta} N_{\geq t,\theta} \underline{X}_{\theta,k}$. If $X_{t-1,\theta,k}^{alg} = \overline{X}_{\theta,k}$, then

$$B_{t,k}^{alg} = B_{t-1,k}^{alg} - \sum_{\theta \in \Theta} N_{t-1,\theta} X_{t-1,\theta,k}^{alg}$$

$$\overset{(a)}{\geq} \sum_{\theta \in \Theta} \underline{X}_{\theta,k}(\mathbb{E}[N_{\geq t,\theta}] + \textsc{Conf}_{t-1,\theta}) \overset{(b)}{\geq} \sum_{\theta \in \Theta} N_{\geq t,\theta} \underline{X}_{\theta,k}.$$

Here, (a) holds by the condition for allocating $\overline{X}_{\theta,k}$, and (b) holds under event $\mathcal{E}$. □

The next lemma shows that the algorithm is adaptively cautious, that is, after some point, Guarded-Hope switches to allocating according to the lower threshold.

Lemma 4. *For each resource k, let $t_{0,k}$ be the last time that $X_{t,\theta,k}^{alg} \neq \overline{X}_{\theta,k}$ (or else zero if the algorithm always allocates according to $\overline{X}_{\theta,k}$). Then, under the event $\mathcal{E}$ for some $c = \tilde{\Theta}(1)$, we have that, for all k, $t_{0,k} \geq \max\{0, T - 2cL_T^{-2}\}$.*

Proof of Lemma 4. We show the result by contradiction (for the setting when $T - 2cL_T^{-2} > 0$). For some resource k, assume that $0 < t_{0,k} < T - 2cL_T^{-2}$. By definition of $t_{0,k}$, it must be that the algorithm allocated $\underline{X}_{\theta,k}$ at time $t_{0,k}$ and allocated $\overline{X}_{\theta,k}$ for all subsequent times. Given the assumption, it must be that, for any $t > t_{0,k}$,

$$\begin{aligned}
&\sum_\theta \underline{X}_{\theta,k} \; (\mathbb{E}[N_{>t,\theta}] + \textsc{Conf}_{t,\theta}) \\
&\overset{(a)}{\leq} B_{t,k}^{alg} - \sum_\theta N_{t,\theta} \overline{X}_{\theta,k} \\
&\overset{(b)}{=} B_{t_{0,k}}^{alg} - \sum_\theta \underline{X}_{\theta,k} N_{t_0,\theta} - \sum_\theta \overline{X}_{\theta,k} \sum_{i=t_{0,k}+1}^{t} N_{i,\theta} \\
&\overset{(c)}{<} \sum_\theta \underline{X}_{\theta,k}(\mathbb{E}[N_{>t_{0,k},\theta}] + \textsc{Conf}_{t_{0,k},\theta}) \\
&\quad + \sum_\theta (\overline{X}_{\theta,k} - \underline{X}_{\theta,k}) N_{t_{0,k},\theta} - \sum_\theta \overline{X}_{\theta,k} \sum_{i=t_{0,k}+1}^{t} N_{i,\theta},
\end{aligned}$$

where (a) follows from the condition in the algorithm for $X_{t,\theta,k}^{alg} = \overline{X}_{\theta,k}$, (b) follows from the definition of $t_{0,k}$ and the choice of $t > t_{0,k}$, and (c) follows from the condition in the algorithm for $X_{t_{0,k},\theta,k}^{alg} = \underline{X}_{\theta,k}$. By rearranging the inequality, we get that

$$\begin{aligned}
&\sum_\theta \overline{X}_{\theta,k} N_{(t_{0,k},t],\theta} \\
&\quad < \sum_\theta (\overline{X}_{\theta,k} - \underline{X}_{\theta,k}) N_{t_{0,k},\theta} + \sum_\theta \underline{X}_{\theta,k} \mathbb{E}[N_{(t_{0,k},t],\theta}] \\
&\qquad + \sum_\theta \underline{X}_{\theta,k}(\textsc{Conf}_{t_{0,k},\theta} - \textsc{Conf}_{t,\theta}) \\
&\quad \Leftrightarrow \sum_\theta \overline{X}_{\theta,k}(N_{(t_{0,k},t],\theta} - \mathbb{E}[N_{(t_0,t],\theta}]) \\
&\quad < \sum_\theta (\overline{X}_{\theta,k} - \underline{X}_{\theta,k})(N_{t_{0,k},\theta} - \mathbb{E}[N_{(t_{0,k},t],\theta}]) \\
&\qquad + \sum_\theta \underline{X}_{\theta,k}(\textsc{Conf}_{t_{0,k},\theta} - \textsc{Conf}_{t,\theta}).
\end{aligned}$$

Using the fact that $\overline{X}_{\theta,k} - \underline{X}_{\theta,k} \leq L_T \frac{\|B\|_\infty \|\beta_{avg}\|_{min} \|w\|_{min}}{\|w\|_\infty}$, and plugging in an upper bound for the demand at time $t_{0,k}$, a lower bound on the expected demand, and the confidence terms from Lemma EC.3, the right-hand side of the inequality can be bounded above by

$$\begin{aligned}
&L_T \frac{\|B\|_\infty \|\beta_{avg}\|_{min} \|w\|_{min}}{\|w\|_\infty} |\Theta| (\rho_{max} + \mu_{max}) \\
&\quad - L_T \frac{\|B\|_\infty \|\beta_{avg}\|_{min} \|w\|_{min}}{\|w\|_\infty} |\Theta| (t - t_{0,k}) \\
&\quad + \|B\|_\infty |\Theta| \sqrt{2\rho_{max}^2 \log(T|\Theta|/\delta)} (\sqrt{T - t_{0,k}} - \sqrt{T - t}).
\end{aligned}$$

Moreover, the left-hand side can be bounded below under event $\mathcal{E}$ via

$$\begin{aligned}
&\sum_\theta \overline{X}_{\theta,k}(N_{(t_{0,k},t],\theta} - \mathbb{E}[N_{(t_0,t],\theta}]) \\
&\qquad \geq - \|\beta_{avg}\|_\infty \sqrt{2\rho_{max}^2 \log(T\Theta|/\delta)(t - t_{0,k})}.
\end{aligned}$$

Plugging in the value of $t = T - cL_T^{-2}$ and by assumption that $t_{0,k} < T - 2cL_T^{-2}$, it follows that, for $\xi = \sqrt{2\rho_{max}^2 \log(T|\Theta|/\delta)}$, $\zeta = \frac{\|B\|_\infty \|\beta_{avg}\|_{min} \|w\|_{min}}{\|w\|_\infty}$,

$$\begin{aligned}
-\|\beta_{avg}\|_\infty \xi \sqrt{cL_T^{-2}} &\leq L_T |\Theta| \zeta (\rho_{max} + \mu_{max}) - L_T |\Theta| \zeta c L_T^{-2} \\
&\quad + \|B\|_\infty |\Theta| \xi \left(\sqrt{2cL_T^{-2}} - \sqrt{cL_T^{-2}} \right).
\end{aligned}$$

Relabeling $x = \sqrt{c}$ to show a contradiction we need to find a value of x such that $-\frac{a_1}{L_T}x^2 + \frac{a_2}{L_T}x + a_3 L_T \le 0$. Noting that the cusp of the quadratic is at $\frac{a_2}{a_1}$, we see that taking the constant

$$c = \frac{\|B\|_\infty |\Theta| \xi + \|\beta_{avg}\|_\infty \xi}{|\Theta| \zeta}$$

(independent of L_T) suffices to show the contradiction ($\Rightarrow \Leftarrow$). □

The main result follows from these two lemmas. We start by giving the upper bound on $\Delta_{efficiency}$. Following from Lemma 4, which states that, if $t_{0,k}$ denotes the last time that $X^{alg}_{t,\theta,k} \neq \overline{X}_{\theta,k}$, then for all k, $t_{0,k} \ge \max\{0, T - 2cL_T^{-2}\}$ with high probability. This implies that, for all k,

$$\begin{aligned}
\Delta_{efficiency} &= \sum_k \left(B_k - \sum_{t \in [T]} \sum_\theta X^{alg}_{t,\theta,k} N_{i,\theta} \right) \\
&= \sum_k \left(B^{alg}_{t_{0,k}} - \sum_{t \ge t_{0,k}} \sum_{\theta \in \Theta} N_{t,\theta} X^{alg}_{t,\theta,k} \right) \\
&\overset{(a)}{=} \sum_k B^{alg}_{t_{0,k},k} - \sum_\theta \left(\underline{X}_{\theta,k} N_{t_{0,k},\theta} + \sum_{t > t_{0,k}} \overline{X}_{\theta,k} N_{t,\theta} \right) \\
&\overset{(b)}{<} \sum_k \sum_\theta (\underline{X}_{\theta,k} (\mathbb{E}[N_{>t_{0,k},\theta}] + \textsc{Conf}_{t_{0,k},\theta} - N_{>t_{0,k},\theta}) \\
&\quad - (\overline{X}_{\theta,k} - \underline{X}_{\theta,k})(N_{>t_{0,k},\theta} - N_{t_{0,k},\theta})),
\end{aligned}$$

where (a) follows from the fact that, by the definition of $t_{0,k}$, the algorithm allocated the lower allocation at time $t_{0,k}$ and the upper allocation for all $t > t_{0,k}$, and (b) follows from the condition in the algorithm for allocating the lower allocation at time $t_{0,k}$, which upper bounds $B^{alg}_{t_{0,k},k}$.

However, under $\mathcal{E}$, we know that $\mathbb{E}[N_{>t_{0,k},\theta}] - N_{>t_{0,k},\theta} \le 2\,\textsc{Conf}_{t_{0,k},\theta}$. Plugging in the definition of $\textsc{Conf}_{t_{0,k},\theta}$ and the bound on $(\overline{X}_{\theta,k} - \underline{X}_{\theta,k})$ from Theorem 4, we have

$$\begin{aligned}
\Delta_{efficiency} &\le 2 \sum_k \sum_\theta \underline{X}_{\theta,k}\ \textsc{Conf}_{t_{0,k},\theta} + N_{t_{0,k},\theta}(\overline{X}_{\theta,k} - \underline{X}_{\theta,k}) \\
&\le 2\|B\|_1 \sum_k \sum_\theta \sqrt{2\rho^2_{max}(T - t_0)\log(T|\Theta|/\delta)} \\
&\quad + (\mu_{max} + \rho_{max}) \frac{2\|B\|_\infty \|\beta_{avg}\|^2_{min} \|w\|_{min} \max_\theta \|w_\theta\|_1}{\|w\|_\infty} L_T.
\end{aligned}$$

Taking this and plugging in the value of $t_{0,k}$ we get that

$$\begin{aligned}
\Delta_{efficiency} &\le 2\|B\|_1 K|\Theta| \sqrt{2\rho^2_{max}\log(T|\Theta|/\delta)} \min\{\sqrt{T}, \sqrt{2c}/L_T\} \\
&\quad + \frac{2(\mu_{max} + \rho_{max})\|B\|_\infty \|\beta_{avg}\|^2_{min} \|w\|_{min} \max_\theta \|w_\theta\|_1}{\|w\|_\infty} L_T.
\end{aligned}$$

Note that $L_T^2 = o(1)$ such that the second term is dominated by the first.

Next, we show the desired bound on $\textsc{Envy}$. Consider an arbitrary t, θ, t', θ'. Then, we have that

$$\begin{aligned}
&u(X^{alg}_{t',\theta'}, \theta) - u(X^{alg}_{t,\theta}, \theta) \\
&= u(X^{alg}_{t',\theta'}, \theta) - u(\underline{X}_{\theta'}, \theta) + u(\underline{X}_{\theta'}, \theta) - u(\underline{X}_\theta, \theta) \\
&\quad + u(\underline{X}_\theta, \theta) - u(X^{alg}_{t,\theta}, \theta) \\
&\overset{(a)}{\le} u(X^{alg}_{t',\theta'}, \theta) - u(\underline{X}_{\theta'}, \theta) + u(\underline{X}_\theta, \theta) - u(X^{alg}_{t,\theta}, \theta) \\
&\overset{(b)}{\le} \|w_{\theta'}\|_1 \|X^{alg}_{t',\theta'} - \underline{X}_{\theta'}\| + \|w_\theta\|_1 \|X^{alg}_{t,\theta} - \underline{X}_\theta\| \\
&\overset{(c)}{\le} 2 \max_\theta \|w_\theta\|_1 \|\overline{X} - \underline{X}\|_\infty \\
&\overset{(d)}{\le} \frac{2\|B\|_\infty \|\beta_{avg}\|^2_{min} \|w\|_{min} \max_\theta \|w_\theta\|_1}{\|w\|_\infty} L_T,
\end{aligned}$$

where in (a) we use that $\underline{X}$ is envy-free, we know the second pair is bounded above by zero, (b) we use the definition of the utilities and (c) the fact that the algorithm allocates according to guardrails, and (d) the bound in Theorem 4. Taking max over t, t', θ, θ' gives the result.

Next, we show the bound on Δ_{EF}. First, consider the setting when $L_T \ge 2\frac{\|w\|^2_\infty}{\|w\|_{min}\|\beta_{avg}\|_{min}} \sqrt{\frac{2\rho^2_{max}\log(T|\Theta|/\delta)}{T}}$ such that we satisfy properties 4 and 5 of Theorem 4. Using that the algorithm always allocates according to $\overline{X}_\theta$ or $\underline{X}_\theta$ with probability at least $1 - \delta$, we get

$$|u(X^{alg}_{t,\theta}, \theta) - u(X^{opt}_\theta, \theta)| \le |u(\overline{X}_\theta, \theta) - u(\underline{X}_\theta, \theta)| \le L_T.$$

However, even for the case that $L_T < 2\frac{\|w\|^2_\infty}{\|w\|_{min}\|\beta_{avg}\|_{min}} \sqrt{\frac{2\rho^2_{max}\log(T|\Theta|/\delta)}{T}}$, then for any θ, we have either $u(\underline{X}_\theta, \theta) \le u(X^{opt}_\theta, \theta) \le u(\overline{X}_\theta, \theta)$ as in property 5, or $u(\underline{X}_\theta, \theta) \le u(\overline{X}_\theta, \theta) \le u(X^{opt}_\theta, \theta)$. If the first case holds, then

$$|u(X^{alg}_{t,\theta}, \theta) - u(X^{opt}_\theta, \theta)| \le |u(\overline{X}_\theta, \theta) - u(\underline{X}_\theta, \theta)| \le L_T.$$

Otherwise, then we can consider $\tilde{X}_\theta$ to be the upper guardrail solution via the construction from Section 5 with $L_T = 2\frac{\|w\|^2_\infty}{\|w\|_{min}\|\beta_{avg}\|_{min}} \sqrt{\frac{2\rho^2_{max}\log(T|\Theta|/\delta)}{T}}$. There, we have

$$\begin{aligned}
&|u(X^{alg}_{t,\theta}, \theta) - u(X^{opt}_\theta, \theta)| \\
&\quad \le |u(\tilde{X}_\theta, \theta) - u(\underline{X}_\theta, \theta)| \le 2\frac{\|w\|^2_\infty}{\|w\|_{min}\|\beta_{avg}\|_{min}} \\
&\quad \sqrt{\frac{2\rho^2_{max}\log(T|\Theta|/\delta)}{T}}.
\end{aligned}$$

Finally, we show the bound on Δ_{prop}. Recall that $\textsc{Guarded-Hope}$ satisfies that $\Delta_{EF} \lesssim \max\{1/\sqrt{T}, L_T\}$. However, by definition of Δ_{EF}, this ensures that, for any round t and type θ that $|u(X^{alg}_{t,\theta},\theta) - u(X^{opt}_{\theta},\theta)| \lesssim \max\{1/\sqrt{T}, L_T\}$. Using this and the fact that X^{opt} is proportional, we see that

$$\begin{aligned} u(\beta_{avg},\theta) - u(X^{alg}_{t,\theta},\theta) &= u(\beta_{avg},\theta) - u(X^{opt}_{\theta},\theta) + u(X^{opt}_{\theta},\theta) - u(X^{alg}_{t,\theta},\theta) \\ &\lesssim \max\{1/\sqrt{T}, L_T\}. \end{aligned}$$

Taking the max over t and θ gives the desired bound on Δ_{prop}. □

7.1. Randomization and the Convex Envelope

As stated in Theorems 2, 3, and 5, the given upper and lower bounds are precise up to $o(1)$ factors. Recall that the performance guarantee on Guarded-Hope with parameter L_T is $\Delta_{efficiency} \lesssim \min\{\sqrt{T}, 1/L_T\}$. We can improve the performance of the algorithm for values of $L_T \in [0, 1/\sqrt{T}]$ by randomization. Indeed, let π_{L_T} denote the Guarded-Hope allocation policy with parameter L_T. Consider $\pi(\alpha)$ to be the allocation policy that picks π_0 with probability $(1-\alpha)$ and $\pi_{1/\sqrt{T}}$ with probability α, playing that policy once chosen across all rounds. It is easy to see that this policy achieves expected metrics

$$\textsc{Envy}(\pi(\alpha)) = \frac{\alpha}{\sqrt{T}} \quad \text{and} \quad \Delta_{efficiency}(\pi(\alpha)) = (1-\alpha)\sqrt{\mathrm{T}} + \frac{\alpha}{\sqrt{\mathrm{T}}}.$$

This improves the performance on $\Delta_{efficiency}$ for values of L_T smaller than $1/\sqrt{T}$.

7.2. Generalizing Distributional Assumptions

Theorem 5 considers the setting when $N_{t,\theta} \sim \mathcal{F}_t$ is independent across θ and a time-dependent process with bounded mean absolute deviation and finite variance. However, it is simple to see that the proofs only require a bound on the following event:

$$\bigcap_{t,\theta} \mathcal{E}_{t,\theta} \quad \text{where}$$

$$\mathcal{E}_{t,\theta} = \{|N_{>t,\theta} - \mathbb{E}[N_{>t,\theta}|N_{\le t,\theta}]| \le \textsc{Conf}_{t,\theta}(N_{\le t,\theta})\}$$

with the scaling of $\textsc{Conf}_{t,\theta}$ being on the order of $\sqrt{T-t}$. The main modification to Guarded-Hope is to condition on the observed sequence of $N_{\le t,\theta}$ thus far, particularly in step two of the algorithm:

(Sufficient budget to promise lower threshold) If $B^{alg}_{t,k} \ge \overline{X}_{\theta,k} N_{t,\theta} + \underline{X}_{\theta,k}(\mathbb{E}[N_{>t,\theta} \mid N_{\le t,\theta}] + \textsc{Conf}_{t,\theta}(N_{\le t,\theta}))$, then set $X^{alg}_{t,\theta,k} = \overline{X}_{\theta,k}\ \forall \theta \in \Theta$.

The concentration arguments and scaling of $\textsc{Conf}_{t,\theta}$ are used in three different sections:

1. Construction of the lower and upper guardrails ($\overline{X}_\theta$ and $\underline{X}_\theta$). This uses concentration of $\mathbb{E}[N_\theta]$, which is given by $\textsc{Conf}_{0,\theta}$.

2. Ensuring the algorithm doesn't run out of budget by saving resources to allocate $\underline{X}_\theta$ to every individual (as in Lemma 3). This utilizes the confidence intervals on $N_{>t,\theta}$, which would be given by $\textsc{Conf}_{t,\theta}(N_{\le t,\theta})$ now conditional on the observed $N_{\le t,\theta}$ values.

3. Construction of the time point after which Guarded-Hope switches to allocating to the lower threshold (as in Lemma 4) uses the scaling of $\textsc{Conf}_{t,\theta}$ as $\sqrt{T-t}$.

Each of these steps is given by $\cap_{t,\theta}\mathcal{E}_{t,\theta}$, and so our approach works under distributional assumptions that yield Chernoff bounds on each event $\mathcal{E}_{t,\theta}$:

- $N_{t,\theta} \sim \mathcal{F}_t$ is a time-dependent process in which each $N_{t,\theta}$ has bounded mean absolute deviation and finite variance (as in Lemma EC.3).
- $N_{t,\theta} \sim \mathcal{F}_t$ is a time-dependent process in which each $N_{t,\theta}$ is sub-Gaussian.
- $N_{t,\theta}$ are conditionally independent and sub-Gaussian given a latent variable Z. This model naturally encompasses dependence on the weather, other local events, etc.
- $N_\theta = \sum_t N_{t,\theta}$ is known for each θ. In this setting, the algorithm can simply take $\mathbb{E}[N_{>t,\theta}|N_{\le t,\theta}] + \textsc{Conf}_{t,\theta}(N_{\le t,\theta}) = N_\theta - N_{\le t,\theta}$.
- Each $N_{t,\theta}$ evolves independently across θ according to different ergodic Markov chains. Concentration bounds for these processes can be constructed using recent work on Chernoff–Hoeffding bounds for Markov chains (theorem 3.1 in Chung et al. 2012).

8. Conclusion

In this paper, we consider the problem of dividing limited resources to individuals arriving over T rounds, and each round can be thought of a distribution location. In the off-line setting (in which the number of individuals arriving to each location is known), achieving a fair allocation scheme is found by maximizing the Nash social welfare objective subject to budget constraints. However, in online settings, no online algorithm can achieve fairness properties ex post. We instead consider the objective of minimizing Δ_{EF} (the maximum difference between the utility individuals receive from the allocation made by the algorithm and the counterfactual optimal fair allocation in hindsight), Envy (the maximum difference between the utility individuals receive from the allocation made by the algorithm and the allocation given to a different individual), and $\Delta_{efficiency}$ (the additive excess of resources).

We show that this objective leads to the envy-efficiency uncertainty principle, an exact characterization between the achievable (Δ_{EF}, Envy, $\Delta_{efficiency}$) pairs. In particular, our result shows that envy and efficiency must be inversely proportional to one another. With this analysis, we show that it leads to a simple algorithm, Guarded-Hope, which is obtained by solving the Eisenberg– Gale program with unknown quantities replaced with their expectation to generate guardrails used in the allocation, combined with an adaptive algorithm aimed at minimizing waste. Through experiments, we show that GUARDED-HOPE is able to obtain allocations that achieve any fairness–efficiency trade-off with desirable fairness properties compared with several benchmarks.

Several open questions remain, including extending the analysis to more general utility functions (including homothetic, another common model of preferences over resources). We also believe much of the theoretical results apply to settings in which the budget B is instead a stochastic process, accounting for external donations and depletions of the resources independent of the allocations made by the algorithm. Moreover, we leave the question of matching the upper and lower bounds in terms of problem dependent constants and the issue of determining the schedule to visit locations as future work.

Acknowledgments

Part of this work was done when Sean Sinclair and Christina Yu were visiting the Simons Institute for the Theory of Computing for the semester on the Theory of Reinforcement Learning. The authors also thank the Food Bank of the Southern Tier and the Cornell Mathematical Contest in Modeling for their collaboration. Finally, the results and their presentation have been significantly improved by the excellent feedback from anonymous reviewers.

References

Aleksandrov M, Walsh T (2019) Monotone and online fair division. *Joint German/Austrian Conf. Artificial Intelligence* (Springer), 60–75.

Aleksandrov M, Walsh T (2020) Online fair division: A survey. *34th AAAI Conf. Artificial Intelligence* (AAAI Press), 13557–13562.

Aleksandrov MD, Aziz H, Gaspers S, Walsh T (2015) Online fair division: Analysing a food bank problem. *24th Internat. Joint Conf. Artificial Intelligence*.

Alkaabneh F, Diabat A, Gao HO (2020) A unified framework for efficient, effective, and fair resource allocation by food banks using an approximate dynamic programming approach. *Omega* 100:102300.

Azar Y, Buchbinder N, Jain K (2010) How to allocate goods in an online market? *Eur. Sympos. Algorithms* (Springer), 51–62.

Aziz H, Schlotter I, Walsh T (2016) Control of fair division. Kambhampati S, ed. *Proc. 25th Internat. Joint Conf. Artificial Intelligence* (IJCAI/AAAI Press), 67–73.

Banerjee S, Gkatzelis V, Gorokh A, Jin B (2020) Online Nash social welfare maximization via promised utilities. Preprint, submitted August 20, https://arxiv.org/abs/2008.03564.

Bansal N, Jiang H, Singla S, Sinha M (2020) Online vector balancing and geometric discrepancy. *Proc. 52nd Annual ACM SIGACT Sympos. Theory Comput.*, 1139–1152.

Bateni M, Chen Y, Ciocan DF, Mirrokni V (2022) Fair resource allocation in a volatile marketplace. *Oper. Res.* 70(1):288–308.

Bauer L (2020) About 14 million children in the US are not getting enough to eat. *Brookings Online* (July 9), https://www.brookings.edu/blog/up-front/2020/07/09/about-14-million-children-in-the-us-are-not-getting-enough-to-eat/.

Benade G, Kazachkov AM, Procaccia AD, Psomas CA (2018) How to make envy vanish over time. *Proc. 2018 ACM Conf. Econom. Comput.*, 593–610.

Berry AC (1941) The accuracy of the Gaussian approximation to the sum of independent variates. *Trans. Amer. Math. Soc.* 49(1):122–136.

Bertsekas D (2012) *Dynamic Programming and Optimal Control*, vol. 1 (Athena Scientific, Nashua, NH).

Bogomolnaia A, Moulin H (2001) A new solution to the random assignment problem. *J. Econom. Theory* 100(2):295–328.

Bogomolnaia A, Moulin H, Sandomirskiy F (2022) On the fair division of a random object. *Management Sci.* 68(2):1174–1194.

Brams SJ, Taylor AD (1995) An envy-free cake division protocol. *Amer. Math. Monthly* 102(1):9–18.

Brams SJ, Taylor AD (1996) *Fair Division: From Cake-Cutting to Dispute Resolution* (Cambridge University Press).

Chung KM, Lam H, Liu Z, Mitzenmacher M (2012) Chernoff-Hoeffding bounds for Markov chains: Generalized and simplified. *29th Internat. Sympos. Theoret. Aspects Comput. Sci.*

Cole R, Gkatzelis V, Goel G (2013) Mechanism design for fair division: Allocating divisible items without payments. *Proc. 14th ACM Conf. Electronic Commerce*, 251–268.

Devanur NR, Papadimitriou CH, Saberi A, Vazirani VV (2002) Market equilibrium via a primal-dual-type algorithm. *Proc. 43rd Annual IEEE Sympos. Foundations Comput. Sci.* (IEEE), 389–395.

Donahue K, Kleinberg J (2020) Fairness and utilization in allocating resources with uncertain demand. *Proc. 2020 Conf. Fairness Accountability Transparency*, 658–668.

Eisenberg E (1961) Aggregation of utility functions. *Management Sci.* 7(4):337–350.

Eisenhandler O, Tzur M (2019) The humanitarian pickup and distribution problem. *Oper. Res.* 67(1):10–32.

Elzayn H, Jabbari S, Jung C, Kearns M, Neel S, Roth A, Schutzman Z (2019) Fair algorithms for learning in allocation problems. *Proc. Conf. Fairness Accountability Transparency*, 170–179.

Friedman E, Psomas CA, Vardi S (2015) Dynamic fair division with minimal disruptions. *Proc. 16th ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 697–713.

Friedman E, Psomas CA, Vardi S (2017) Controlled dynamic fair division. *Proc. 2017 ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 461–478.

Gerding EH, Perez-Diaz A, Aziz H, Gaspers S, Marcu A, Mattei N, Walsh T (2019) Fair online allocation of perishable goods and its application to electric vehicle charging. *Proc. 28th Internat. Joint Conf. Artificial Intelligence* (International Joint Conferences on Artificial Intelligence Organization), 5569–5575.

Ghodsi A, Zaharia M, Hindman B, Konwinski A, Shenker S, Stoica I (2011) Dominant resource fairness: Fair allocation of multiple resource types. *NSDI*, vol. 11, 24.

He J, Procaccia AD, Psomas A, Zeng D (2019) Achieving a fairer future by changing the past. *Proc. 28th Internat. Joint Conf. Artificial Intelligence* (International Joint Conferences on Artificial Intelligence Organization), 343–349.

Jaberi-Douraki M, Moghadas SM (2014) Optimal control of vaccination dynamics during an influenza epidemic. *Math. Biosciences Engrg.* 11(5):1045–1063.

Kalinowski T, Narodytska N, Walsh T (2013) A social welfare optimal sequential allocation procedure. *23rd Internat. Joint Conf. Artificial Intelligence*.

Kash I, Procaccia AD, Shah N (2014) No agent left behind: Dynamic fair division of multiple resources. *J. Artificial Intelligence Res.* 51:579–603.

Kleinberg J, Mullainathan S, Raghavan M (2016) Inherent trade-offs in the fair determination of risk scores. Preprint, submitted September 19, https://arxiv.org/abs/1609.05807.

Kulish N (2020) "Never seen anything like it": Cars line up for miles at food banks. *The New York Times Online* (April 8), https://www.nytimes.com/2020/04/08/business/economy/coronavirus-food-banks.html.

Lien RW, Iravani SM, Smilowitz KR (2014) Sequential resource allocation for nonprofit operations. *Oper. Res.* 62(2):301–317.

Lupkin S (2020) How feds decide on remdesivir shipments to states remains mysterious. *NPR Online* (August 19), https://www.npr.org/sections/health-shots/2020/08/19/903946857/how-feds-decide-on-remdesivir-shipments-to-states-remains-mysterious.

Manshadi V, Niazadeh R, Rodilitz S (2021) Fair dynamic rationing. Preprint, submitted February 2, https://dx.doi.org/10.2139/ssrn.3775895.

Mattei N, Saffidine A, Walsh T (2017) Mechanisms for online organ matching. *Proc. 26th Internat. Joint Conf. Artificial Intelligence*, 345–351.

Mattei N, Saffidine A, Walsh T (2018) Fairness in deceased organ matching. *Proc. 2018 AAAI/ACM Conf. AI Ethics Society*, 236–242.

Moulin H (2004) *Fair Division and Collective Welfare* (MIT Press, Cambridge, MA).

Nisan N, Roughgarden T, Tardos É, Vazirani VV, eds. (2007) *Algorithmic Game Theory* (Cambridge University Press, Cambridge, UK).

Orgut IS, Brock LG III, Davis LB, Ivy JS, Jiang S, Morgan SD, Uzsoy R, Hale C, Middleton E (2016) Achieving equity, effectiveness, and efficiency in food bank operations: Strategies for feeding America with implications for global hunger relief. *Advances in Managing Humanitarian Operations* (Springer), 229–256.

Procaccia AD (2013) Cake cutting: Not just child's play. *Comm. ACM.* 56(7):78–87.

Sengul Orgut I, Ivy J, Uzsoy R (2017) Modeling for the equitable and effective distribution of food donations under stochastic receiving capacities. *IISE Trans.* 49(6):567–578.

Shadmi E, Chen Y, Dourado I, Faran-Perach I, Furler J, Hangoma P, Hanvoravongchai P, et al. (2020) Health equity and COVID-19: Global perspectives. *Internat. J. Equity Health* 19(1):1–16.

Sugden R (1984) Is fairness good? A critique of Varian's theory of fairness. *Nous* 18(3):505–511.

Varian HR (1974) Equity, envy, and efficiency. *J. Econom. Theory* 9(1):63–91.

Varian HR (1976) Two problems in the theory of fairness. *J. Public Econom.* 5(3–4):249–260.

Vera A, Banerjee S (2019) The Bayesian prophet: A low-regret framework for online decision making. *Performance Evaluation Rev.* 47(1):81–82.

Walsh T (2011) Online cake cutting. *Internat. Conf. Algorithmic Decision Theory* (Springer), 292–305.

Yi M, Marathe A (2015) Fairness vs. efficiency of vaccine allocation strategies. *Value Health* 18(2):278–283.

Zeng D, Psomas A (2020) Fairness-efficiency tradeoffs in dynamic fair division. *Proc. 21st ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 911–912.

Sean R. Sinclair is a graduate student in the school of operations research and information engineering at Cornell University, working on developing algorithms for data-driven sequential decision making in societal applications.

Gauri Jain is a graduate student in computer science at Harvard University.

Siddhartha Banerjee is an associate professor in the school of operations research and information engineering at Cornell University, working on topics at the intersection of data-driven decision making, network algorithms, and market design.

Christina Lee Yu is an assistant professor at Cornell University in the school of operations research and information engineering.