Appendix A: Additional Related Work

This section surveys additional works beyond those covered in Section 2 that are also related to this work.

Beyond worst-case Analysis for OLP: Beyond worst-case approaches for OLP problems have focused on designing algorithms under (i) the random permutation and (ii) the stochastic input models. In the random permutation model, the constraints and objective coefficients arrive according to a random permutation of an adversarially chosen input sequence. In this context, Devanur and Hayes (2009), Agrawal et al. (2014) develop a two-phase algorithm, which includes training the model on a small fraction of the input sequence and then using the learned parameters to make online decisions on the remaining input sequence. Contrastingly, in the stochastic input model, the input sequence is drawn i.i.d. from some potentially unknown distribution. In this setting, Li et al. (2022) investigate the convergence of the dual price vector and design algorithms using LP duality to obtain logarithmic regret bounds. Since the algorithms developed in Li et al. (2022) involve solving an LP at specified intervals, Li et al. (2022), Gao et al. (2024) developed gradient descent-based algorithms wherein the dual prices are adjusted solely based on the allocation to users at each time step. Furthermore, Chen et al. (2024) devised an adaptive allocation algorithm with constant regret when the samples are drawn from a discrete distribution. As with some of these works, we develop algorithms for online Fisher markets under the stochastic input model; however, in contrast to these works that assume a linear objective, we develop regret guarantees for a non-linear concave objective function.

Revealed Preferences: Our approach of adjusting prices using users' revealed preferences, i.e., observed user consumption information, is analogous to price update mechanisms that use information from interactions with earlier buyers to inform pricing decisions for future buyers (Kleinberg and Leighton 2003). While our dual-based price update mechanism is akin to those used in prior work on revealed preferences (Roth et al. 2016, Ji et al. 2018), our work considers a setting with budget-constrained users, unlike the quasi-linear utility setting studied in these works. Prior literature on revealed preference has also considered the setting of budget-constrained users (Zadimoghaddam and Roth 2012, Bei et al. 2016, Balcan et al. 2014, Beigman and Vohra 2006) as in this work. However, these works focus on the problem of learning the budgets and valuation functions of users that rationalize their observed buying behavior rather than designing algorithms with low regret, which is one of the main focuses of this work.

Online Constraint Convex Optimization with Long-term Constraints: Furthermore, since we focus on jointly optimizing regret and constraint violation, our work closely relates to the literature on online constrained convex optimization with long-term constraints (Yi et al. 2021, Liakopoulos et al. 2019, Jenatton et al. 2016, Mahdavi et al. 2012, Valls et al. 2020). However, compared to these works that focus on a regret measure defined based on the sub-optimality of an optimal static action in hindsight, we adopt a more powerful oracle model, wherein the oracle can vary its actions across time steps as in Yu et al. (2017), Chen et al. (2017), Cao and Liu (2019). Even though our chosen regret metric is akin to the dynamic regret notions in these works, our work differs from Yu et al. (2017), Chen et al. (2017), Cao and Liu (2019) in several ways. First, unlike these works, which consider a setting wherein the central planner observes a convex cost function after each user arrival, we study a revealed preference setting, wherein users' utility and budget parameters are private information. Next, as opposed to the gradient descent projection step used in the algorithms developed in Yu et al. (2017), Chen et al. (2017), Cao and Liu (2019), we establish regret and constraint violation bounds for our revealed preference algorithm without projecting the price vector to the non-negative orthant. We do so by developing a novel potential function argument that relies on the structural properties of Fisher markets (see Section 5.3 and Appendix F). Finally, compared to Yu et al. (2017), Chen et al. (2017), Cao and Liu (2019), we also consider the informational setting when the distribution $\mathcal{D}$ is discrete and known to the central planner and develop an adaptive expected equilibrium pricing algorithm in this setting with constant constraint violation and logarithmic regret.

Artificial Currency Mechanisms: Our work is also closely related to the design and analysis of artificial currency mechanisms (Kash et al. 2007, Gorokh et al. 2021). Such mechanisms have found applications in various resource allocation settings, including the allocation of food to food banks (Prendergast 2016), the allocation of students to courses (Budish 2011), and the allocation of public goods to people (Jalota et al. 2020). Mechanisms that involve artificial currencies have also been designed for repeated allocation settings (Gorokh et al. 2016), as is the main focus of this paper. However, unlike Gorokh et al. (2016) that studies the repeated allocation of goods that arrive online, we investigate the setting of online user arrival.

Appendix B: Regret and Nash Social Welfare

We establish a fundamental connection between the regret measure studied in this work and the ratio between the Nash social welfare objective of the optimum offline oracle and that corresponding to an online algorithm. In particular, we show that if the regret $U_n^* - U_n(\boldsymbol{\pi}) \leq o(n)$ for some algorithm $\boldsymbol{\pi}$, then $\frac{NSW(\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)}{NSW(\mathbf{x}_1, \ldots, \mathbf{x}_n)} \to 1$ as $n \to \infty$. Here, $\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*$ are the optimal offline allocations, and $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are the optimal consumption vectors given by the solution of Problem (1a)-(1c) under the prices corresponding to the online pricing policy $\boldsymbol{\pi}$. Without loss of generality, consider the setting when the budgets of all users are equal. Note that if the budgets are not equal, then we can just re-scale the utilities of each user based on their budget. In this setting, it holds that

$$\frac{1}{n}U_n^* = \frac{1}{n}\sum_{t=1}^n \log(u_t(\mathbf{x}_t^*)) = \frac{1}{n}\log\left(\prod_{t=1}^n u_t(\mathbf{x}_t^*)\right) = \log\left(\left(\prod_{t=1}^n u_t(\mathbf{x}_t^*)\right)^{\frac{1}{n}}\right) = \log(NSW(\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)),$$

and $\frac{1}{n}U_n(\boldsymbol{\pi}) = \log\left(\left(\prod_{t=1}^n u_t(\mathbf{x}_t)\right)^{\frac{1}{n}}\right) = \log(NSW(\mathbf{x}_1, \ldots, \mathbf{x}_n))$. Then, it follows that

$$\frac{NSW(\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)}{NSW(\mathbf{x}_1, \ldots, \mathbf{x}_n)} = \frac{e^{\frac{1}{n}U_n^*}}{e^{\frac{1}{n}U_n(\boldsymbol{\pi})}} = e^{\frac{1}{n}(U_n^* - U_n(\boldsymbol{\pi}))} \leq e^{\frac{o(n)}{n}}.$$

Observe that as $n \to \infty$, the term $e^{\frac{o(n)}{n}} \to 1$. That is, if the regret of an algorithm $\boldsymbol{\pi}$ is $o(n)$, then the ratio of the Nash social welfare of algorithm $\boldsymbol{\pi}$ approaches that of the optimal offline oracle as n becomes large.

Appendix C: Proof of Theorem 1

Consider a setting with n users with a fixed budget of one and two goods, each with a capacity of n. Further, let the utility parameters of users be drawn i.i.d. from a distribution, where the users have utility $(1,0)$ with probability 0.5 and a utility of $(0,1)$ with probability 0.5. That is, users only have utility for good one or good two, each with equal probability. For this instance, we first derive a tight bound for the expected optimal social welfare objective, i.e., Objective (2a). Then, to establish the desired lower bound, we consider two cases: (i) the price of either of the two goods is at most 0.5, and (ii) the price of both goods is strictly greater than 0.5. In the first case, we establish that the expected constraint violation is $\Omega(\sqrt{n})$ while in the second case, we establish that either the expected constraint violation or the expected regret is $\Omega(\sqrt{n})$.

C.1. Tight Bound on Expected Optimal Social Welfare Objective

To obtain a bound on the expected optimal social welfare objective, we first find an expression for the objective given the number of arrivals s of users with the utility $(1,0)$. To this end, for the defined problem instance, given s arrivals of users with the utility $(1,0)$ (for ease of exposition, let the first s indexed users have a utility of $(1,0)$), we have the following offline social optimization problem

$$U^*(s) = \max_{\substack{\mathbf{x}_t \in \mathbb{R}^2, \\ \forall t \in [n]}} \quad \sum_{t=1}^{s} \log(x_{t1}) + \sum_{t=s+1}^{n} \log(x_{t2}), \tag{6a}$$

$$\text{s.t.} \quad \sum_{t=1}^{n} x_{t1} \leq n, \tag{6b}$$

$$\sum_{t=1}^{n} x_{t2} \leq n, \tag{6c}$$

$$x_{tj} \geq 0, \quad \forall t \in [n], j \in [2]. \tag{6d}$$

If $0 < s < n$, then the optimal solution of the above problem is to allocate $\mathbf{x}_t = (\frac{n}{s}, 0)$ to each user t with a utility of $(1,0)$ and to allocate $\mathbf{x}_t = (0, \frac{n}{n-s})$ to each user t with a utility of $(0,1)$. In this case, the optimal objective value is given by

$$U^*(s) = s \log\left(\frac{n}{s}\right) + (n-s)\log\left(\frac{n}{n-s}\right) = n\log(n) - s\log(s) - (n-s)\log(n-s).$$

We now develop a tight bound on the expected optimal objective $U^*(s)$ using the fact that the number of arrivals s of users with utility $(0,1)$ is binomially distributed with a probability of 0.5. That is, we seek to develop a tight bound for

$$\begin{aligned}\mathbb{E}[U^*(s)] &= \mathbb{E}[n\log(n) - s\log(s) - (n-s)\log(n-s)], \\ &= n\log(n) - \mathbb{E}[s\log(s)] - \mathbb{E}[(n-s)\log(n-s)].\end{aligned}$$

To this end, we present an upper bound for $s\log(s)$ and $(n-s)\log(n-s)$, which will yield a lower bound for $\mathbb{E}[U^*(s)]$.

We begin by observing that the expectation of the binomial random variable is given by $\mathbb{E}[s] = \frac{n}{2}$ and its variance is $\mathbb{E}[(s-\frac{n}{2})^2] = \frac{n}{4}$. Next, letting $\sigma = \frac{2}{n}(s-\frac{n}{2})$, which has zero mean and a standard deviation of $\frac{1}{\sqrt{n}}$, we obtain the following upper bound on the term $s\log(s)$

$$
\begin{aligned}
s\log(s) &= s\log\left(\frac{n}{2}+s-\frac{n}{2}\right), \\
&= s\log\left(\frac{n}{2}\left(1+\frac{2}{n}\left(s-\frac{n}{2}\right)\right)\right), \\
&= s\log\left(\frac{n}{2}\right)+s\log\left(1+\frac{2}{n}\left(s-\frac{n}{2}\right)\right), \\
&= s\log\left(\frac{n}{2}\right)+s\log(1+\sigma), \\
&\leq s\log\left(\frac{n}{2}\right)+s\sigma. \qquad (7)
\end{aligned}
$$

Similarly, we obtain the following upper bound for $(n-s)\log(n-s)$:

$$
(n-s)\log(n-s) = (n-s)\log\left(\frac{n}{2}\right)+(n-s)\log(1-\sigma) \leq (n-s)\log\left(\frac{n}{2}\right)-(n-s)\sigma \qquad (8)
$$

Adding Equations (7) and (8), we have that

$$
s\log(s)+(n-s)\log(n-s) \leq n\log\left(\frac{n}{2}\right)+(2s-n)\sigma = n\log\left(\frac{n}{2}\right)+n\sigma^2.
$$

As a result, it holds that

$$
U^*(s) = n\log(n)-s\log(s)-(n-s)\log(n-s) \geq n\log(n)-n\log\left(\frac{n}{2}\right)-n\sigma^2 = n\log(2)-n\sigma^2 \qquad (9)
$$

for all $0<s<n$. Next, letting q_s be the probability of observing s users with utility $(1,0)$, it follows that

$$
\begin{aligned}
\mathbb{E}[U^*(s)] &= \sum_{s=0}^{n} q_s U^*(s) \overset{(a)}{=} \sum_{s=1}^{n-1} q_s U^*(s) \overset{(b)}{\geq} \sum_{s=1}^{n-1} q_s (n\log(2)-n\sigma^2), \\
&\overset{(c)}{\geq} \left(1-\frac{1}{2^{n-1}}\right) n\log(2)-n\mathbb{E}[\sigma^2], \\
&= \left(1-\frac{1}{2^{n-1}}\right) n\log(2)-n\mathbb{E}\left[\left(\frac{2}{n}\left(s-\frac{n}{2}\right)\right)^2\right], \\
&= \left(1-\frac{1}{2^{n-1}}\right) n\log(2)-1
\end{aligned}
$$

where (a) follows as $U^*(0)=0$ and $U^*(n)=0$, (b) follows by Equation (9), (c) follows as $\sum_{s=1}^{n-1} q_s = 1-\frac{1}{2^{n-1}}$ and $\sum_{s=1}^{n-1} q_s\sigma^2 \leq \sum_{s=0}^{n} q_s\sigma^2 = \mathbb{E}[\sigma^2]$.

Finally, using Jensen's inequality for a concave function, we obtain the following upper bound on the expected optimal social welfare objective:

$$
\mathbb{E}[U^*(s)] \leq U^*(\mathbb{E}(s)) \leq n\log(2).
$$

As a result, we have shown the following tight bound on the expected optimal social welfare objective for the earlier defined instance:

$$
\left(1-\frac{1}{2^{n-1}}\right) n\log(2)-1 \leq \mathbb{E}[U^*(s)] \leq n\log(2).
$$

C.2. Lower bound on Expected Regret and Constraint Violation

Case (i): We first consider the case when the price of either of the two goods is at most 0.5. Without loss of generality, let $p_1 \leq 0.5$. Then, with s arrivals of users with utility $(1,0)$, the expected constraint violation of good one is given by

$$v_1 = \mathbb{E}\left[\left(\frac{s}{p_1} - n\right)_+\right] \geq \mathbb{E}\left[(2s-n)_+\right],$$

which is $O(\sqrt{n})$ by the central limit theorem as $\frac{n}{2}$ users of each type arrive in expectation. As a result, the norm of the constraint violation $\Omega(\sqrt{n})$. This establishes that if the price of either of the goods is below 0.5, the expected constraint violation is $\Omega(\sqrt{n})$.

Case (ii): Next, we consider the case when the price of both goods is strictly greater than 0.5. In particular, suppose that $\mathbf{p} = (p_1, p_2) = (\frac{1}{2-\epsilon_1(n)}, \frac{1}{2-\epsilon_2(n)})$, where $\epsilon_1(n), \epsilon_2(n) > 0$ can depend on the number of users n and are constants for any fixed value of n. We now show that for any choice of $\epsilon_1(n), \epsilon_2(n) > 0$ that either the expected regret or the expected constraint violation is $\Omega(\sqrt{n})$.

To this end, first note by the central limit theorem that the expected constraint violation for good one for s arrivals of users with utility $(1,0)$ is given by

$$v_1 = \mathbb{E}\left[\left(\frac{s}{p_1} - n\right)_+\right] = \mathbb{E}\left[(s(2-\epsilon_1(n)) - n)_+\right] \geq \mathbb{E}\left[(2s-n)_+\right] - \epsilon_1(n)\mathbb{E}[s] = \Omega(\sqrt{n}) - \epsilon_1(n)\frac{n}{2}. \quad (10)$$

Similarly, the expected constraint violation of good two is lower bounded by $\Omega(\sqrt{n}) - \epsilon_2(n)\frac{n}{2}$.

Next, using the lower bound on the expected optimal social welfare objective we obtain the following lower bound on the regret of any static pricing policy with $\mathbf{p} = (\frac{1}{2-\epsilon_1(n)}, \frac{1}{2-\epsilon_2(n)})$:

$$\begin{aligned}
\text{Regret} &\geq \left(1 - \frac{1}{2^{n-1}}\right) n\log(2) - 1 - \mathbb{E}\left[\sum_{t=1}^n \log\left(\frac{1}{\frac{1}{2-\epsilon(n)}}\right)\right], \\
&= \left(1 - \frac{1}{2^{n-1}}\right) n\log(2) - 1 - \mathbb{E}\left[\sum_{t=1}^n \log(2-\epsilon(n))\right], \\
&= \left(1 - \frac{1}{2^{n-1}}\right) n\log(2) - 1 - \mathbb{E}\left[\sum_{t=1}^n \log\left(2\left(1 - \frac{\epsilon(n)}{2}\right)\right)\right], \\
&= -\frac{1}{2^{n-1}} n\log(2) - 1 - \mathbb{E}\left[\sum_{t=1}^n \log\left(1 - \frac{\epsilon(n)}{2}\right)\right], \\
&\geq -\frac{1}{2^{n-1}} n\log(2) - 1 + \frac{n\epsilon(n)}{2},
\end{aligned}$$

where $\epsilon(n) = \min\{\epsilon_1(n), \epsilon_2(n)\}$ and $0 < \epsilon(n) < 2$.

Finally, to simultaneously achieve the lowest regret and constraint violation, we set $\Omega(\sqrt{n}) - \epsilon(n)\frac{n}{2} = -1 + \frac{n\epsilon(n)}{2} - \frac{1}{2^{n-1}} n\log(2)$. Solving for $\epsilon(n)$, we get that $\epsilon(n) = O(\frac{1}{\sqrt{n}})$ as n becomes large. This relation implies that to minimize both regret and constraint violation, $\epsilon(n)$ needs to be set on the order of $\frac{1}{\sqrt{n}}$, which will result in a corresponding expected regret and constraint violation of $\Omega(\sqrt{n})$. Observe that for any other choice of $\epsilon(n)$, either the regret or the constraint violation must be $\Omega(\sqrt{n})$ since setting $\epsilon(n) = O(\frac{1}{\sqrt{n}})$ guarantees that both the regret and constraint violation are minimized. This establishes our claim that either the regret or the constraint violation must be $\Omega(\sqrt{n})$ when the price of both goods is strictly greater than 0.5, which proves our claim.

Appendix D: Proof of Theorem 2

We prove Theorem 2 using four intermediate lemmas, which we elucidate below. After presenting the statements of these lemmas, we then present their proofs.

Our first lemma establishes a generic upper bound on the regret of an algorithm for the online Fisher market setting considered in this work. To define this generic regret bound, we first introduce the following stochastic program

$$\min_{\mathbf{p}} \quad D(\mathbf{p}) = \sum_{j=1}^{m} p_j d_j + \mathbb{E}\left[\left(w \log(w) - w \log\left(\min_{j \in [m]} \frac{p_j}{u_j}\right) - w\right)\right], \tag{11}$$

which is the stochastic programming formulation of the dual of the Eisenberg-Gale program (see Equation (5)) presented in Section 5.1. Letting $\mathbf{p}^*$ be the optimal solution to this stochastic program, we obtain the following generic bound on the regret of any algorithm for online Fisher markets.

LEMMA 1 **(Generic Regret Bound).** *Suppose that the budget and utility parameters of users are drawn i.i.d. from a probability distribution $\mathcal{D}$. Furthermore, let $\boldsymbol{\pi}$ denote an online pricing policy, $\mathbf{x}_1, \ldots, \mathbf{x}_n$ be the corresponding allocations for the n users, and $\underline{p}, \bar{p} > 0$ be the lower and upper bounds, respectively, for the prices p_j^t for all goods j and for all users $t \in [n]$, where the price upper bound $\bar{p} \geq \max_{j \in [m]} p_j^*$. Then, the regret $R_n(\boldsymbol{\pi}) \leq \frac{2\sqrt{m}\bar{w}}{\underline{p}} \sum_{t=1}^{n} \mathbb{E}\left[\|\mathbf{p}^* - \mathbf{p}^t\|_2\right] + \mathbb{E}\left[\bar{p}\left|\sum_{j=1}^{m} \left(\sum_{t=1}^{n} x_{tj} - c_j\right)\right|\right]$.*

A few comments about Lemma 1 are in order. First, observe that the generic regret bound obtained in Lemma 1 applies to general (non-discrete) probability distributions $\mathcal{D}$. Next, the generic regret bound is composed of two terms: (i) the first term accounts for the loss from setting prices that deviate from the optimal expected prices $\mathbf{p}^*$, and (ii) the second term is akin to the constraint violation of the algorithm and, in particular, accounts for the loss corresponding to over (or under-consuming) certain goods.

As a result, to upper bound the regret of Algorithm 1, we now present lemmas that upper bound both the terms in the generic regret upper bound. To this, end, we first show that the upper bound on the expected constraint violation is constant in the number of arriving users. This result not only establishes the desired constraint violation bound in the statement of Theorem 2 but its analysis also provides a bound on the second term of the generic regret upper bound in Lemma 1.

LEMMA 2 **(Constraint Violation Bound of Algorithm 1).** *Suppose that the budget and utility parameters of users are drawn i.i.d. from a discrete probability distribution $\mathcal{D}$ and let $\boldsymbol{\pi}$ denote the online pricing policy* *described by Algorithm 1. Furthermore, let $\mathbf{x}_1, \dots, \mathbf{x}_n$ be the corresponding allocations for the n users, where $\mathbf{x}_t$ is an optimal solution for that user corresponding to the certainty equivalent problem $CE(\mathbf{d}_t)$ for $t \leq \tau$, where τ is the first time at which $\mathbf{d}_t \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta]$, and $\mathbf{x}_t$ is an optimal solution to $CE(\mathbf{d})$ for $t > \tau$. Then, the constraint violation $V_n(\boldsymbol{\pi}) \leq O(1)$.*

The proof of Lemma 2 follows from an application of similar techniques to that used in Chen et al. (2024). In this proof, we leverage the fact that the allocations $\mathbf{x}_t$ are given by the optimal solution of the certainty equivalent problem $CE(\mathbf{d}_t)$ for $t \leq \tau$, which is one of the optimal consumption vectors corresponding to the price $\mathbf{p}^t$. Note that doing so is without loss of generality, since the utility of the users is unchanged for any optimal consumption bundle. Furthermore, recall from Section 4.2.1 that the allocations corresponding to the optimal solution of the certainty equivalent problem $CE(\mathbf{d}_t)$ at each step can be implemented in Algorithm 1 using an allocation-based algorithm, wherein users are given allocations based on their observed type $k \in [K]$.

Having obtained a bound on the constraint violation, we next upper bound the first term in the generic regret upper bound. To do so, we proceed in two steps. First, we establish a Lipschitzness relation between the optimal price vector of the certainty equivalent problem $CE(\mathbf{d}_t)$ and the average remaining resource capacity vector $\mathbf{d}_t$, as is elucidated through the following lemma.

LEMMA 3 **(Lipschitz Relation Between Prices and Average Remaining Resource Capacities).** *Suppose $\mathbf{d}, \mathbf{d}' > \mathbf{0}$ are two average remaining resource capacity vectors and $\mathbf{p}^*(\mathbf{d}), \mathbf{p}^*(\mathbf{d}')$ are the optimal price vectors corresponding to the certainty equivalent problems $CE(\mathbf{d}), CE(\mathbf{d}')$, respectively. Then, $\|\mathbf{p}^*(\mathbf{d}) - \mathbf{p}^*(\mathbf{d}')\|_2 \leq L \|\mathbf{d} - \mathbf{d}'\|_2$ for some constant $L > 0$.*

Lemma 3 establishes that small changes in the average remaining capacity vector will only result in small changes in the corresponding optimal price vector of the certainty equivalent Problem $CE(\mathbf{d}_t)$. In particular, Lemma 3 implies that if $\|\mathbf{d} - \mathbf{d}'\|_2 \leq O(\frac{1}{n-t})$ for a given $t \in [n-1]$, then, the optimal price vectors $\mathbf{p}, \mathbf{p}'$ of the certainty equivalent problems $CE(\mathbf{d})$ and $CE(\mathbf{d}')$, respectively, satisfy $\|\mathbf{p} - \mathbf{p}'\|_2 \leq O(\frac{1}{n-t})$. We also numerically validate this obtained Lipschitz relation in Appendix M.1.

We then leverage Lemma 3 to establish an $O(\log(n))$ upper bound on the first term of the generic regret bound, as is elucidated through the following lemma.

LEMMA 4 **(Bound on Difference in Prices).** *Suppose that the budget and utility parameters of users are drawn i.i.d. from a discrete probability distribution $\mathcal{D}$. Furthermore, let $\boldsymbol{\pi}$ denote the online pricing policy described by Algorithm 1 and let $\mathbf{x}_1, \dots, \mathbf{x}_n$ be the corresponding allocations for the n users. Then, $\frac{2\sqrt{m}\bar{w}}{\underline{p}} \sum_{t=1}^n \mathbb{E}\left[\|\mathbf{p}^* - \mathbf{p}^t\|_2\right] \leq O(\log(n))$.*

Finally, we combine the results obtained in Lemmas 2 and 4 to obtain the $O(\log(n))$ upper bound on the regret of Algorithm 1.

COROLLARY 5 **(Regret Upper Bound of Algorithm 1).** *Suppose that the budget and utility parameters of users are drawn i.i.d. from a discrete probability distribution $\mathcal{D}$ and let $\boldsymbol{\pi}$ denote the online pricing policy described by Algorithm 1. Furthermore, let $\mathbf{x}_1, \dots, \mathbf{x}_n$ be the corresponding allocations for the n users, where $\mathbf{x}_t$ is an optimal solution for that user corresponding to the certainty equivalent problem $CE(\mathbf{d}_t)$ for $t \leq \tau$, where τ is the first time at which $\mathbf{d}_t \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta]$, and $\mathbf{x}_t$ is an optimal solution to $CE(\mathbf{d})$ for $t > \tau$. Then, the regret $R_n(\boldsymbol{\pi}) \leq O(\log(n))$.*

Note that Lemma 2 and Corollary 5 jointly imply Theorem 2, which thus proves our claim.

D.1. Proof of Lemma 1

We now establish a generic bound on the regret of any online algorithm as long as the prices $\mathbf{p}^t$ are strictly positive and bounded, i.e., $0 < \underline{p} \leq p_j^t \leq \bar{p}$ for all goods j and for all users $t \in [n]$. To establish a generic upper bound on the regret, we first obtain a bound on the expected value of the optimal objective, i.e., Objective (2a), and a relation for the expected value of the objective for any online allocation policy $\boldsymbol{\pi}$. We finally combine both these relations to obtain an upper bound on the regret.

To perform our analysis, we define the function $g(\mathbf{p}) = \mathbb{E}[w_t \log(\mathbf{u}_t^T \mathbf{x}_t) + \sum_{j=1}^m (d_j - x_{tj}(\mathbf{p}))p_j^*]$, where $\mathbf{p}^*$ is the optimal price vector of the stochastic Program (11). Then, by duality we have that the expected primal objective value $\mathbb{E}[U_n^*]$ is no more than the dual objective value with $\mathbf{p} = \mathbf{p}^*$, which gives the following upper bound on the optimal objective

$$\begin{aligned} \mathbb{E}[U_n^*] &\leq \mathbb{E}\left[\sum_{j=1}^m p_j^* c_j + \sum_{t=1}^n \left(w_t \log(w_t) - w_t \log\left(\min_{j \in [m]} \frac{p_j^*}{u_{tj}}\right) - w_t\right)\right], \\ &= nD(\mathbf{p}^*), \end{aligned} \tag{12}$$

by the definition of $D(\mathbf{p})$ in Problem (11). Next, we establish a relation between the function $g(\mathbf{p})$ and the above obtained bound on the expected value of the optimal objective value by noting that

$$\begin{aligned} g(\mathbf{p}^*) &= \mathbb{E}[w_t \log(\mathbf{u}_t^T \mathbf{x}_t^*) + \sum_{j=1}^m (d_j - x_{tj}^*(\mathbf{p}^*))p_j^*], \\ &\overset{(a)}{=} \mathbb{E}\left[w_t \log(w_t) - w_t \log\left(\min_{j \in [m]} \frac{p_j^*}{u_{tj}}\right) + \sum_{j \in [m]} p_j^* d_j - w_t\right], \\ &= nD(\mathbf{p}^*), \end{aligned} \tag{13}$$

where (a) follows by the definition of $g(\mathbf{p})$ and noting that for each agent $t \in [n]$ it holds that $\mathbf{u}_t^T \mathbf{x}_t = u_{tj'} \frac{w_t}{p_{j'}^*}$ for some good j' in the optimal bundle for the user t, and that $\sum_{j \in [m]} x_{tj}(\mathbf{p}^*)p_j^* = w_t$ since each user spends their entire budget when consuming its optimal bundle of goods given the price vector $\mathbf{p}^*$. Combining the relations obtained in Equations (12) and (13), we obtain the following upper bound on the expected value of the optimal objective:

$$\mathbb{E}[U_n^*] \leq ng(\mathbf{p}^*). \tag{14}$$

Having obtained an upper bound on the expected optimal objective, we now obtain the following relationship for the true accumulated social welfare objective, i.e., Objective (2a), accrued by any online policy $\boldsymbol{\pi}$ that sets prices $\mathbf{p}^1, \dots, \mathbf{p}^n$ with corresponding allocations $\mathbf{x}_1, \dots, \mathbf{x}_n$:

$$\begin{aligned}
\mathbb{E}\left[U_n(\boldsymbol{\pi})\right] &= \mathbb{E}\left[\sum_{t=1}^{n} w_t \log(\mathbf{u}_t^T \mathbf{x}_t)\right], \\
&= \mathbb{E}\left[\sum_{t=1}^{n} w_t \log(\mathbf{u}_t^T \mathbf{x}_t) + \sum_{j=1}^{m} p_j^* \left(c_j - \sum_{t=1}^{n} x_{tj}\right) - \sum_{j=1}^{m} p_j^* \left(c_j - \sum_{t=1}^{n} x_{tj}\right)\right], \\
&= \mathbb{E}\left[\sum_{t=1}^{n}\left(w_t \log(\mathbf{u}_t^T \mathbf{x}_t) + \sum_{j=1}^{m} p_j^* \left(d_j - x_{tj}\right)\right)\right] + \mathbb{E}\left[\sum_{j=1}^{m} p_j^* \left(\sum_{t=1}^{n} x_{tj} - c_j\right)\right] \qquad (15)
\end{aligned}$$

We can analyse the first term on the right hand side of Equation (15) as follows:

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{n}\left(w_t \log(\mathbf{u}_t^T \mathbf{x}_t) + \sum_{j=1}^{m} p_j^* \left(d_j - x_{tj}\right)\right)\right] &\overset{(a)}{=} \sum_{t=1}^{n} \mathbb{E}\left[w_t \log(\mathbf{u}_t^T \mathbf{x}_t) + \sum_{j=1}^{m} p_j^* \left(d_j - x_{tj}\right)\right], \\
&\overset{(b)}{=} \sum_{t=1}^{n} \mathbb{E}\left[\mathbb{E}\left[w_t \log(\mathbf{u}_t^T \mathbf{x}_t) + \sum_{j=1}^{m} p_j^* \left(d_j - x_{tj}\right) | \mathcal{H}_{t-1}\right]\right], \\
&\overset{(c)}{=} \sum_{t=1}^{n} \mathbb{E}\left[g(\mathbf{p}^t)\right] = \mathbb{E}\left[\sum_{t=1}^{n} g(\mathbf{p}^t)\right], \qquad (16)
\end{aligned}$$

where (a) follows by the linearity of expectation, (b) follows from nesting conditional expectations, where the history $\mathcal{H}_{t-1} = \{w_i, \mathbf{u}_i, \mathbf{x}_i\}_{i=1}^{t-1}$, and (c) follows from the definition of $g(\mathbf{p})$ and the fact that the allocation x_{tj} depends on the vector of prices $\mathbf{p}_t$.

Finally, combining the above analysis in Equations (14), (15), and (16) for $\mathbb{E}[U_n^*]$ and $\mathbb{E}[U_n(\boldsymbol{\pi})]$, we obtain the following bound on the regret of any online allocation policy $\boldsymbol{\pi}$ for $\bar{p} \geq \max_{j \in [m]} p_j^*$:

$$\begin{aligned}
\mathbb{E}[U_n^* - U_n(\boldsymbol{\pi})] &\leq n g(\mathbf{p}^*) - \mathbb{E}\left[\sum_{t=1}^{n} g(\mathbf{p}^t)\right] - \mathbb{E}\left[\sum_{j=1}^{m} p_j^* \left(\sum_{t=1}^{n} x_{tj} - c_j\right)\right], \\
&\leq \mathbb{E}\left[\sum_{t=1}^{n}(g(\mathbf{p}^*) - g(\mathbf{p}^t))\right] + \mathbb{E}\left[\bar{p}\left|\sum_{j=1}^{m}\left(\sum_{t=1}^{n} x_{tj} - c_j\right)\right|\right]. \qquad (17)
\end{aligned}$$

Finally, to obtain the desired generic regret bound, we establish that $\mathbb{E}\left[g(\mathbf{p}^*) - g(\mathbf{p}^t)\right] \leq O(\mathbb{E}\left[\|\mathbf{p}^* - \mathbf{p}^t\|_2\right])$. To this end, first observe from the definition of the function g that for the optimal solution $\mathbf{x}_t(\mathbf{p})$ of the individual optimization Problem (1a)-(1c) given a price vector $\mathbf{p}$ that

$$\begin{aligned}
g(\mathbf{p}^*) - g(\mathbf{p}^t) &= \mathbb{E}\left[w_t \log(\mathbf{u}_t^T \mathbf{x}_t(\mathbf{p}^*)) + \sum_{j=1}^{m}(d_j - x_{tj}(\mathbf{p}^*)) p_j^*\right] \\
&\quad - \mathbb{E}\left[w_t \log(\mathbf{u}_t^T \mathbf{x}_t(\mathbf{p}^t)) + \sum_{j=1}^{m}(d_j - x_{tj}(\mathbf{p}_t)) p_j^*\right], \\
&= \mathbb{E}\left[w_t \log\left(\min_{j \in [m]}\left\{\frac{p_j^t}{u_{tj}}\right\} \frac{1}{\min_{j \in [m]}\left\{\frac{p_j^*}{u_{tj}}\right\}}\right)\right] + \mathbb{E}\left[\sum_{j=1}^{m}(x_{tj}(\mathbf{p}^t) - x_{tj}(\mathbf{p}^*)) p_j^*\right].
\end{aligned}$$

Then, letting the good $j' \in \arg\min_{j \in [m]} \{\frac{p_j^*}{u_{tj}}\}$ and $j^*(\mathbf{p})$ be a good in the optimal consumption set of user t given the price $\mathbf{p}$, we observe that

$$
\begin{aligned}
g(\mathbf{p}^*) - g(\mathbf{p}_t) &\overset{(a)}{\leq} \mathbb{E}\left[w_t \log\left(\frac{p_{j'}^t}{p_{j'}^*}\right)\right] + \mathbb{E}\left[\sum_{j=1}^m \left(\mathbb{1}_{j=j^*(\mathbf{p}^t)} \frac{w_t}{p_j^t} - \mathbb{1}_{j=j^*(\mathbf{p}^*)} \frac{w_t}{p_j^*}\right) p_j^*\right], \\
&\overset{(b)}{=} \mathbb{E}\left[w_t \log\left(1 + \frac{p_{j'}^t - p_{j'}^*}{p_{j'}^*}\right)\right] + \mathbb{E}\left[\sum_{j=1}^m \frac{w_t(p_j^* - p_j^t)}{p_j^* p_j^t} \left(\mathbb{1}_{j=j^*(\mathbf{p}^t)} - \mathbb{1}_{j=j^*(\mathbf{p}^*)}\right) p_j^*\right], \\
&\overset{(c)}{\leq} \mathbb{E}\left[w_t \frac{p_{j'}^t - p_{j'}^*}{p_{j'}^*}\right] + \mathbb{E}\left[\sum_{j=1}^m \frac{w_t(p_j^* - p_j^t)}{p_j^t}\right], \\
&\overset{(d)}{\leq} \frac{2\bar{w}}{\underline{p}} \mathbb{E}\left[\left\|\mathbf{p}^* - \mathbf{p}^t\right\|_1\right], \\
&\overset{(e)}{\leq} \frac{2\sqrt{m}\bar{w}}{\underline{p}} \mathbb{E}\left[\left\|\mathbf{p}^* - \mathbf{p}^t\right\|_2\right],
\end{aligned} \tag{18}
$$

where (a) follows since $j' \in \arg\min_{j \in [m]} \{\frac{p_j^*}{u_{tj}}\}$ and $\mathbf{x}_t(\mathbf{p})$ corresponds to the optimal solution to the individual optimization Problem (1a)-(1c), (b) follows by rearranging the right hand side of the equation in (a). Next, (c) follows from the fact that $\log(1+x) \leq x$ for $x > -1$ and that the difference between two indicators can be at most one. Inequality (d) follows by the upper bound on the budgets of users and the lower bound on the price vector. The final inequality (e) follows from the norm equivalence property which holds for the one and two norms.

Finally, using Equations (18) and (17), we obtain the following generic upper bound on the regret of any online algorithm $\boldsymbol{\pi}$:

$$
\mathbb{E}[U_n^* - U_n(\boldsymbol{\pi})] \leq \frac{2\sqrt{m}\bar{w}}{\underline{p}} \sum_{t=1}^n \mathbb{E}\left[\left\|\mathbf{p}^* - \mathbf{p}^t\right\|_2\right] + \mathbb{E}\left[\bar{p}\left|\sum_{j=1}^m \left(\sum_{t=1}^n x_{tj} - c_j\right)\right|\right], \tag{19}
$$

which proves our claim.

D.2. Proof of Lemma 2

To prove this result, we first prove an upper bound on the expected constraint violation in terms of the stopping time τ of the algorithm. Then, we establish a lower bound on the expected value of the stopping time to establish the constant constraint violation bound.

Upper Bound on constraint violation in terms of stopping time: We begin by establishing that the constraint violation of Algorithm 1 is upper bounded by $O(\mathbb{E}[n-\tau])$, where the stopping time $\tau = \min\{t \leq n : \mathbf{d}_t \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta]\} \cup \{n\}$. To this end, first note by the definition of τ and that $\Delta < \mathbf{d}$ that no constraints are violated up until user τ. Furthermore, since the consumption $x_{tj} \leq \frac{\bar{w}}{\underline{p}}$ for all $t > \tau$, it follows that the constraint violation

$$
\mathbb{E}\left[\left\|\sum_{t=1}^n x_{tj} - c_j\right\|_2\right] \leq \mathbb{E}\left[\left\|\sum_{t=\tau+1}^n x_{tj}\right\|_2\right] \leq \mathbb{E}\left[(n-\tau)\sqrt{m}\frac{\bar{w}}{\underline{p}}\right] = O(\mathbb{E}[n-\tau]). \tag{20}
$$

Bound on Expected Stopping time τ: From the above analysis, we observe that bounding the expected constraint violation amounts to obtaining a bound on the expected stopping time τ. To this end, we first introduce some notation. In particular, as in Chen et al. (2024), we define the following auxiliary process:

$$\tilde{\mathbf{d}}_t = \begin{cases} \mathbf{d}_t, & t < \tau \\ \mathbf{d}_\tau, & t \geq \tau \end{cases}.$$

Then, we can obtain a generic bound on the expected stopping time by observing that

$$\begin{aligned}
\mathbb{E}[\tau] &= \sum_{t=1}^n t\mathbb{P}(\tau = t) = \sum_{t=1}^n \mathbb{P}(\tau \geq t) = \sum_{t=1}^n (1 - \mathbb{P}(\tau < t)) \geq \sum_{t=1}^n (1 - \mathbb{P}(\tau \leq t)), \\
&\overset{(a)}{=} \sum_{t=1}^n \left[1 - \mathbb{P}(\mathbf{d}_s \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta] \text{ for some } s \leq t)\right], \\
&\overset{(b)}{\geq} n - \sum_{t=1}^n \mathbb{P}\left(\tilde{\mathbf{d}}_s \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta] \text{ for some } s \leq t\right),
\end{aligned}$$

where (a) follows by the definition of τ, (b) follows since the auxiliary process $\tilde{\mathbf{d}}_s$ is identical to $\mathbf{d}_s$ for all s less than τ. The above analysis implies that

$$\mathbb{E}[n - \tau] \leq \sum_{t=1}^n \mathbb{P}\left(\tilde{\mathbf{d}}_s \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta] \text{ for some } s \leq t\right). \tag{21}$$

Thus, to obtain an upper bound for $\mathbb{E}[n - \tau]$, we now proceed to finding an upper bound for the term $\mathbb{P}\left(\tilde{\mathbf{d}}_s \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta] \text{ for some } s \leq t\right)$ for each user $t \in [n]$.

Upper bound on $\mathbb{P}\left(\tilde{\mathbf{d}}_s \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta] \text{ for some } s \leq t\right)$: To obtain an upper bound on this term, we leverage Hoeffding's inequality:

LEMMA 5. *(Hoeffding's Inequality (van de Geer 2002)) Suppose there is a sequence of random variables. $\{X_t\}_{t=1}^n$ adapted to a filtration $\mathcal{H}_{t-1}$, and $\mathbb{E}[X_t|\mathcal{H}_{t-1}] = 0$ for all $t \in [n]$, where $\mathcal{H}_0 = \emptyset$. Suppose further that L_t and U_t are $\mathcal{H}_{t-1}$ measurable random variables such that $L_t \leq X_t \leq U_t$ almost surely for all $t \in [n]$. Then, letting $S_t = \sum_{s=1}^t X_t$ and $V_t = \sum_{s=1}^t (U_s - L_s)^2$, the following inequality holds for any constants $b, c > 0$: $\mathbb{P}(|S_t| \geq b, V_t \leq c^2 \text{ for some } t \in \{1, \ldots, T\}) \leq 2e^{-\frac{2b^2}{c^2}}$.*

To leverage Lemma 5, we begin by introducing some notation. First define $Y_{tj} := \tilde{d}_{j,t+1} - \tilde{d}_{j,t}$ and $X_{tj} := Y_{tj} - \mathbb{E}[Y_{tj} \mid \mathcal{H}_{t-1}]$ for $t \geq 1$, where $\mathcal{H}_{t-1} = ((w_1, \mathbf{u}_1), \ldots, (w_{t-1}, \mathbf{u}_{t-1}))$ is the history of observed budget and utility parameters.

Next, observe for $t \geq \tau$ that $\tilde{d}_{j,t+1} = \tilde{d}_{j,t}$ and when $1 \leq t < \tau$ we have that:

$$\tilde{d}_{j,t+1} = d_{j,t+1} = \frac{c_{j,t+1}}{n-t} = \frac{c_{jt} - x_{tj}}{n-t} = d_{jt} - \frac{1}{n-t}(x_{tj} - d_{jt}) = \tilde{d}_{jt} - \frac{1}{n-t}(x_{tj} - \tilde{d}_{jt})$$

Next, noting that $\tilde{d}_{jt}$ is $\mathcal{H}_{t-1}$ measurable, we have that:

$$\begin{aligned}
|X_{tj}| &= \left| \frac{1}{n-t}(x_{tj} - \tilde{d}_{jt}) - \mathbb{E}\left[\frac{1}{n-t}(x_{tj} - \tilde{d}_{jt})|\mathcal{H}_{t-1}\right] \right| \\
&= \frac{1}{n-t}|x_{tj} - \mathbb{E}[x_{tj}|\mathcal{H}_{t-1}]| \leq \frac{\bar{w}}{\underline{p}(n-t)}
\end{aligned}$$

for each $t \leq n-1$ due to the boundedness of the allocations x_{tj}. Then, defining $L_t = -\frac{\bar{w}}{\underline{p}(n-t)}$ and $U_t = \frac{\bar{w}}{\underline{p}(n-t)}$, we obtain that

$$V_t = \sum_{s=1}^{t} (U_s - L_s)^2 = \sum_{s=1}^{t} \frac{4\bar{w}^2}{\underline{p}^2 (n-s)^2} \leq \frac{4\bar{w}^2}{\underline{p}^2 (n-t-1)},$$

which holds for all $t \leq n-2$.

Then, from a direct application of Hoeffding's inequality (Lemma 5) for some constant $\Delta' > 0$ we have that

$$\mathbb{P}\left(\left| \sum_{i=1}^{s} X_{ij} \right| \geq \Delta' \text{ for some } s \leq t \right) \leq 2e^{-\frac{\underline{p}^2 \Delta'^2 (n-t-1)}{2\bar{w}^2}}. \tag{22}$$

Next, we observe that

$$|X_{tj} - Y_{tj}| = |\mathbb{E}\left[Y_{tj} \mid \mathcal{H}_{t-1}\right]| = \left| \mathbb{E}\left[\tilde{d}_{j,t+1} - \tilde{d}_{j,t} \mid \mathcal{H}_{t-1} \right] \right| \stackrel{(a)}{=} \left| \frac{1}{n-t} \mathbb{E}\left[\left(x_{tj} - \tilde{d}_{j,t} \right) I(t < \tau) \mid \mathcal{H}_{t-1} \right] \right| = 0, \tag{23}$$

where (a) follows since the probability distribution is exactly known in Algorithm 1 and at the optimal solution of the certainty equivalent problem it holds that $\mathbb{E}[\mathbf{x}_t] = \mathbf{d}_t$ for all $t \leq \tau$. Consequently, it holds that the term $\mathbb{E}\left[\left(x_{tj} - \tilde{d}_{j,t} \right) I(t < \tau) \mid \mathcal{H}_{t-1} \right] = 0$ for all users $t < \tau$.

Then, to obtain a bound on $\mathbb{P}(\tilde{\mathbf{d}}_s \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta]$ for some $s \leq t)$, we first note the following key relation for the set $\{\tilde{\mathbf{d}}_s \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta]$ for some $s \leq t\}$:

$$\begin{aligned}
\left\{ \left| \tilde{d}_{j,s} - d_j \right| > \Delta_j \text{ for some } s \leq t \right\} &\stackrel{(a)}{=} \left\{ \left| \sum_{i=1}^{s-1} Y_{ij} \right| > \Delta_j \text{ for some } s \leq t \right\}, \\
&= \left\{ \left| \sum_{i=1}^{s} Y_{ij} \right| > \Delta_j \text{ for some } s \leq t-1 \right\}, \\
&\stackrel{(b)}{=} \left\{ \left| \sum_{i=1}^{s} X_{ij} \right| > \Delta_j \text{ for some } s \leq t-1 \right\},
\end{aligned}$$

where (a) follows from the definition of Y_i, and (b) follows since $\sum_{i=1}^{s} X_{ij} = \sum_{i=1}^{s} Y_{ij}$, as proved in Equation (23). Then setting $\Delta' = \min_{j \in [m]} \Delta_j = \underline{\Delta}$ in Equation (22), and applying a union bound over all the goods $j \in [m]$, we obtain that

$$\mathbb{P}(\tilde{\mathbf{d}}_s \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta] \text{ for some } s \leq t) \leq 2me^{-\frac{\underline{p}^2 \underline{\Delta}^2 (n-t-1)}{2\bar{w}^2}}, \tag{24}$$

which holds for all $t \leq n-2$.

Constant Bound on Expected Constraint Violation: We have already observed from our earlier analysis that the expected constraint violation is upper bounded by $O(\mathbb{E}[n-\tau])$, where

$$\mathbb{E}[n-\tau] \leq \sum_{t=1}^{n} \mathbb{P}\left(\tilde{\mathbf{d}}_s \notin [\mathbf{d} - \Delta, \mathbf{d} + \Delta] \text{ for some } s \leq t \right)$$

follows from Equation (21). Thus, we now use the obtained upper bound on $\mathbb{P}(\tilde{\mathbf{d}}_s \notin [\mathbf{d}-\Delta, \mathbf{d}+\Delta]$ for some $s \leq t)$ (Equation (24)) for any $t \leq n-2$ to show that $\mathbb{E}[n-\tau]$ is bounded above by a constant. To see this, observe that

$$\begin{aligned}
\mathbb{E}[n-\tau] &\leq 2 + \sum_{t=1}^{n-2} \mathbb{P}(\tilde{\mathbf{d}}_s \notin [\mathbf{d}-\Delta, \mathbf{d}+\Delta] \text{ for some } s \leq t), \\
&\leq 2 + \sum_{t=1}^{n-2} 2me^{-\frac{\underline{p}^2\Delta^2(n-t-1)}{2\bar{w}^2}}, \\
&= 2 + 2m \sum_{s=1}^{n-2} e^{-\frac{\underline{p}^2\Delta^2 s}{2\bar{w}^2}}, \\
&= 2 + 2me^{-\frac{\underline{p}^2\Delta^2}{2\bar{w}^2}} \sum_{s=0}^{n-3} e^{-\frac{\underline{p}^2\Delta^2 s}{2\bar{w}^2}}, \\
&= 2 + 2me^{-\frac{\underline{p}^2\Delta^2}{2\bar{w}^2}} \frac{1-e^{-\frac{\underline{p}^2\Delta^2(n-3)}{2\bar{w}^2}}}{1-e^{-\frac{\underline{p}^2\Delta^2}{2\bar{w}^2}}}, \\
&\leq 2 + 2m \frac{1}{1-e^{-\frac{\underline{p}^2\Delta^2}{2\bar{w}^2}}}, \\
&= O(m)
\end{aligned} \tag{25}$$

The above analysis for the upper bound on the term $\mathbb{E}[n-\tau]$ along with Equation (20) establishes the constant upper bound on the expected constraint violation for Algorithm 1, as

$$\mathbb{E}\left[\left\|\sum_{t=1}^{n} x_{tj} - c_j\right\|_2\right] \leq O(\mathbb{E}[n-\tau]) \leq O(m).$$

This completes the proof of our claim that the constraint violation of Algorithm 1 is bounded by a constant independent of the number of users n.

D.3. Proof of Lemma 3

We begin by presenting the dual function $D(\mathbf{p}, \mathbf{d})$, which is a function of the price vector $\mathbf{p}$ and parametrized by the per-user resource vector $\mathbf{d}$. In particular, the dual function is represented as

$$D(\mathbf{p}, \mathbf{d}) = \sum_{j=1}^{m} p_j d_j - \sum_{k=1}^{K} q_k \tilde{w}_k \log\left(\min_{j\in[m]} \frac{p_j}{\tilde{u}_{kj}}\right) \tag{26}$$

Note that we have dropped the constant terms independent of the price vector $\mathbf{p}$ from the dual objective (see Section 5.1). We now re-parametrize the dual function with a variable $\alpha_k = \min_{j\in[m]} \frac{p_j}{\tilde{u}_{kj}}$ to get the following:

$$D(\boldsymbol{\alpha}, \mathbf{d}) = \sum_{j=1}^{m} d_j \max_k\{\alpha_k \tilde{u}_{kj}\} - \sum_{k=1}^{K} q_k \tilde{w}_k \log(\alpha_k),$$

where note that $\alpha_k \leq \frac{p_j}{\tilde{u}_{kj}}$ for all j and k.

Next, let $\mathbf{d}, \mathbf{d}'$ be two different resource consumption vectors and let $\boldsymbol{\alpha}^*(\mathbf{d})$ (and $\boldsymbol{\alpha}^*(\mathbf{d}')$) be the optimal solution to the corresponding dual problems with resource vectors $\mathbf{d}$ (and $\mathbf{d}'$), respectively. Then, noting that the dual function $D(\boldsymbol{\alpha}, \mathbf{d})$ is strongly convex in $\boldsymbol{\alpha}$ (for a bounded set of values of $\boldsymbol{\alpha}$), it holds that:

$$D(\boldsymbol{\alpha}^*(\mathbf{d}'), \mathbf{d}) - D(\boldsymbol{\alpha}^*(\mathbf{d}), \mathbf{d}) \geq \eta \left\| \boldsymbol{\alpha}^*(\mathbf{d}') - \boldsymbol{\alpha}^*(\mathbf{d}) \right\|^2, \tag{27}$$

where note that $\eta > 0$ is a positive constant as the prices $\mathbf{p}$ remain bounded for all resource consumption vectors $\tilde{\mathbf{d}} \in [\mathbf{d} - \Delta, \mathbf{d} + \Delta]$.

Next, note that

$$\begin{aligned}
D(\boldsymbol{\alpha}^*(\mathbf{d}'), \mathbf{d}) - D(\boldsymbol{\alpha}^*(\mathbf{d}), \mathbf{d}) &= [D(\boldsymbol{\alpha}^*(\mathbf{d}'), \mathbf{d}) - D(\boldsymbol{\alpha}^*(\mathbf{d}'), \mathbf{d}')] - [D(\boldsymbol{\alpha}^*(\mathbf{d}), \mathbf{d}) - D(\boldsymbol{\alpha}^*(\mathbf{d}), \mathbf{d}')] \\
&\quad + D(\boldsymbol{\alpha}^*(\mathbf{d}'), \mathbf{d}') - D(\boldsymbol{\alpha}^*(\mathbf{d}), \mathbf{d}'), \\
&\overset{(a)}{\leq} [D(\boldsymbol{\alpha}^*(\mathbf{d}'), \mathbf{d}) - D(\boldsymbol{\alpha}^*(\mathbf{d}'), \mathbf{d}')] - [D(\boldsymbol{\alpha}^*(\mathbf{d}), \mathbf{d}) - D(\boldsymbol{\alpha}^*(\mathbf{d}), \mathbf{d}')], \\
&= \sum_{j=1}^m d_j \max_k \{\alpha_k^*(\mathbf{d}') \tilde{u}_{kj}\} - \sum_{k=1}^K q_k \tilde{w}_k \log(\alpha_k^*(\mathbf{d}')) \\
&\quad - \sum_{j=1}^m d_j' \max_k \{\alpha_k^*(\mathbf{d}') \tilde{u}_{kj}\} + \sum_{k=1}^K q_k \tilde{w}_k \log(\alpha_k^*(\mathbf{d}')) \\
&\quad - \sum_{j=1}^m d_j \max_k \{\alpha_k^*(\mathbf{d}) \tilde{u}_{kj}\} + \sum_{k=1}^K q_k \tilde{w}_k \log(\alpha_k^*(\mathbf{d})) \\
&\quad + \sum_{j=1}^m d_j' \max_k \{\alpha_k^*(\mathbf{d}) \tilde{u}_{kj}\} - \sum_{k=1}^K q_k \tilde{w}_k \log(\alpha_k^*(\mathbf{d})), \\
&= \sum_{j=1}^m \max_k \{\alpha_k^*(\mathbf{d}') \tilde{u}_{kj}\} (d_j - d_j') + \sum_{j=1}^m \max_k \{\alpha_k^*(\mathbf{d}) \tilde{u}_{kj}\} (d_j' - d_j), \\
&= \sum_{j=1}^m (d_j - d_j') (\max_k \{\alpha_k^*(\mathbf{d}') \tilde{u}_{kj}\} - \max_k \{\alpha_k^*(\mathbf{d}) \tilde{u}_{kj}\}), \\
&\overset{(b)}{\leq} \sum_{j=1}^m (d_j - d_j') (\alpha_{\tilde{k}}^*(\mathbf{d}') \tilde{u}_{\tilde{k}j} - \alpha_{\tilde{k}}^*(\mathbf{d}) \tilde{u}_{\tilde{k}j}), \\
&\leq \bar{u} \left\| \boldsymbol{\alpha}^*(\mathbf{d}) - \alpha^*(\mathbf{d}') \right\|_2 \left\| \mathbf{d} - \mathbf{d}' \right\|_1, \\
&\overset{(c)}{\leq} \sqrt{m} \bar{u} \left\| \boldsymbol{\alpha}^*(\mathbf{d}) - \alpha^*(\mathbf{d}') \right\|_2 \left\| \mathbf{d} - \mathbf{d}' \right\|_2,
\end{aligned}$$

where (a) follows as $\boldsymbol{\alpha}^*(\mathbf{d}')$ is a minimizer of the dual function $D(\cdot, \mathbf{d}')$, (b) follows as we define $\tilde{k}$ to be the group such that $\tilde{k} \in \arg\max\{\alpha_k^*(\mathbf{d}) \tilde{u}_{kj}\}$ and (c) follows by the norm equivalence relation between the one and two norm.

From the above inequality and Equation (27), it follows that

$$\eta \left\| \boldsymbol{\alpha}^*(\mathbf{d}') - \boldsymbol{\alpha}^*(\mathbf{d}) \right\|_2^2 \leq D(\boldsymbol{\alpha}^*(\mathbf{d}'), \mathbf{d}) - D(\boldsymbol{\alpha}^*(\mathbf{d}), \mathbf{d}) \leq \sqrt{m} \bar{u} \left\| \boldsymbol{\alpha}^*(\mathbf{d}) - \alpha^*(\mathbf{d}') \right\|_2 \left\| \mathbf{d} - \mathbf{d}' \right\|_2.$$

Thus, it holds that

$$\|\boldsymbol{\alpha}^*(\mathbf{d}') - \boldsymbol{\alpha}^*(\mathbf{d})\|_2 \leq \frac{\sqrt{m}\bar{u}}{\eta} \|\mathbf{d} - \mathbf{d}'\|_2 . \tag{28}$$

Finally, to establish our desired Lipschitzness result, we show that the prices are Lipschitz in α. To show this, we denote $\mathbf{p}^*(\mathbf{d})$ (and $\mathbf{p}^*(\mathbf{d}')$) be the optimal price vector corresponding to two different average resource consumption vectors $\mathbf{d}$ (and $\mathbf{d}'$), respectively. Then, we get that:

$$\begin{aligned}
\|\mathbf{p}^*(\mathbf{d}) - \mathbf{p}^*(\mathbf{d}')\|_2 &= \left\| \max_k \{\alpha_k^*(\mathbf{d})\mathbf{u}_k\} - \max_k \{\alpha_k^*(\mathbf{d}')\mathbf{u}_k\} \right\|_2 , \\
&= \sqrt{\sum_{j=1}^m \left(\max_k \{\alpha_k^*(\mathbf{d}) u_{kj}\} - \max_k \{\alpha_k^*(\mathbf{d}') u_{kj}\} \right)^2 }, \\
&\overset{(a)}{\leq} \sqrt{\sum_{j=1}^m \left(\alpha_{\tilde{k}_j}^*(\mathbf{d}) u_{\tilde{k}_j j} - \alpha_{\tilde{k}_j}^*(\mathbf{d}') u_{\tilde{k}_j j} \right)^2 }, \\
&\leq \sqrt{m}\bar{u} \|\boldsymbol{\alpha}^*(\mathbf{d}) - \boldsymbol{\alpha}^*(\mathbf{d}')\|_2 ,
\end{aligned} \tag{29}$$

where (a) follows for some $\tilde{k}_j$. To see this, let $\tilde{k}_j^1 = \arg\max_k \{\alpha_k^*(\mathbf{d}) u_{kj}\}$ and let $\tilde{k}_j^2 = \max_k \{\alpha_k^*(\mathbf{d}') u_{kj}\}$. Then, it follows that: $\max_k \{\alpha_k^*(\mathbf{d}) u_{kj}\} - \max_k \{\alpha_k^*(\mathbf{d}') u_{kj}\} \leq \max\{\alpha_{\tilde{k}_j^1}^*(\mathbf{d}) u_{\tilde{k}_j^1 j} - \alpha_{\tilde{k}_j^2}^*(\mathbf{d}') u_{\tilde{k}_j^2 j}, \alpha_{\tilde{k}_j^2}^*(\mathbf{d}') u_{\tilde{k}_j^2 j} - \alpha_{\tilde{k}_j^1}^*(\mathbf{d}) u_{\tilde{k}_j^1 j}\}$ for all goods j.

Finally, combining the above relation with Equation (28), we obtain the desired Lipschitzness relation between prices and the average resource consumption vectors:

$$\|\mathbf{p}^*(\mathbf{d}) - \mathbf{p}^*(\mathbf{d}')\|_2 \leq \sqrt{m}\bar{u} \|\boldsymbol{\alpha}^*(\mathbf{d}) - \boldsymbol{\alpha}^*(\mathbf{d}')\|_2 \leq \frac{m\bar{u}^2}{\eta} \|\mathbf{d} - \mathbf{d}'\|_2 ,$$

which establishes our claim. We also note that the above Lipschitzness relation implies that when $\|\mathbf{d} - \mathbf{d}'\| = O(\frac{1}{n-t})$, then even $\|\mathbf{p}^*(\mathbf{d}') - \mathbf{p}^*(\mathbf{d})\| = O(\frac{1}{n-t})$.

D.4. Proof of Lemma 4

We use Lemma 3 to analyse the first term in the generic regret bound in Equation (19) for Algorithm 1 and establish that $\frac{2\sqrt{m}\bar{w}}{\underline{p}} \sum_{t=1}^n \mathbb{E}\left[\|\mathbf{p}^* - \mathbf{p}^t\|_2\right] \leq O(\log(n))$.

To this end, we first show that $\mathbb{E}[\|\mathbf{p}^t - \mathbf{p}^*\|_2] \leq O(\frac{1}{n-t+1})$ for all $t = \{1, \ldots, n-\tau\}$ for Algorithm 1. To see this, we proceed by induction. For the base case, take $t = 1$, in which case Algorithm 1 initializes the price $\mathbf{p}^1 = \mathbf{p}^*$, as the adaptive expected equilibrium pricing algorithm sets the static expected equilibrium prices at $t = 1$ as $\mathbf{d}_1 = \frac{\mathbf{c}}{n}$. As a result, it clearly holds that $\mathbb{E}[\|\mathbf{p}^1 - \mathbf{p}^*\|_2] = 0 \leq O(\frac{1}{n})$. For the inductive step, we now assume that $\mathbb{E}[\|\mathbf{p}^t - \mathbf{p}^*\|_2] \leq O(\frac{1}{n-t+1})$ for all $t \leq k$. Then, we have for $t = k+1$ that

$$\begin{aligned}
\mathbb{E}\left[\left\|\mathbf{p}^{k+1} - \mathbf{p}^*\right\|_2\right] &\leq \mathbb{E}\left[\left\|\mathbf{p}^{k+1} - \mathbf{p}^k\right\|_2\right] + \mathbb{E}\left[\left\|\mathbf{p}^k - \mathbf{p}^*\right\|_2\right], \\
&\leq \mathbb{E}\left[\left\|\mathbf{p}^{k+1} - \mathbf{p}^k\right\|_2\right] + O\left(\frac{1}{n-k+1}\right), \\
&\leq \mathbb{E}\left[\left\|\mathbf{p}^{k+1} - \mathbf{p}^k\right\|_2\right] + O\left(\frac{1}{n-k}\right).
\end{aligned} \tag{30}$$

To bound $\mathbb{E}[\left\|\mathbf{p}^{k+1}-\mathbf{p}^k\right\|_2]$, we note that $\mathbf{d}_{k+1}=\mathbf{d}_k+\frac{\mathbf{d}_k-\mathbf{x}_k(\mathbf{p}^k)}{n-k}$, i.e., $\|\mathbf{d}_{k+1}-\mathbf{d}_k\|=O(\frac{1}{n-k})$. Then, using Lemma 3, it follows that $\mathbb{E}[\left\|\mathbf{p}^{k+1}-\mathbf{p}^k\right\|_2]=O(\frac{1}{n-k})$. This inequality, together with Equation (30), implies that

$$\mathbb{E}\left[\left\|\mathbf{p}^{k+1}-\mathbf{p}^*\right\|_2\right]\leq O\left(\frac{1}{n-k}\right), \tag{31}$$

which establishes our inductive step and thus establishes our claim that $\mathbb{E}[\|\mathbf{p}^t-\mathbf{p}^*\|_2]\leq O(\frac{1}{n-t+1})$ for all $t=\{1,\dots,n-\tau\}$ for Algorithm 1. Furthermore, observe that since $\mathbf{p}^t=\mathbf{p}^*$ for $t>\tau$, it holds that $\mathbb{E}[\|\mathbf{p}^t-\mathbf{p}^*\|_2]\leq O(\frac{1}{n-t+1})$ for all $t=\{1,\dots,n\}$. Using this result, we obtain the following upper bound on the first term of Equation (19)

$$\frac{2\sqrt{m}\bar{w}}{\underline{p}}\sum_{t=1}^{n}\mathbb{E}\left[\left\|\mathbf{p}^*-\mathbf{p}^t\right\|_2\right]\leq\frac{2\sqrt{m}\bar{w}}{\underline{p}}\sum_{t=1}^{n}\mathbb{E}\left[O\left(\frac{1}{n-t+1}\right)\right]\leq\frac{2\sqrt{m}\bar{w}}{\underline{p}}\sum_{t=1}^{n}\mathbb{E}\left[O\left(\frac{1}{t}\right)\right]\leq O(\log(n)), \tag{32}$$

which proves our claim.

D.5. Proof of Corollary 5

We now use the generic bound on the regret derived in Equation (19) to obtain a $O(\log(n))$ bound on the regret of Algorithm 1. In particular, we upper bound both the terms on the right hand side of Equation (19) using the analysis performed in Lemmas 2 and 4 to establish that

$$\mathbb{E}[U_n^*-U_n(\boldsymbol{\pi})]\leq\frac{2\sqrt{m}\bar{w}}{\underline{p}}\sum_{t=1}^{n}\mathbb{E}\left[\left\|\mathbf{p}^*-\mathbf{p}^t\right\|_2\right]+\mathbb{E}\left[\bar{p}\left|\sum_{j=1}^{m}\left(\sum_{t=1}^{n}x_{tj}-c_j\right)\right|\right]\leq O(\log(n)).$$

To establish the above claim, we first observe by Lemma 4 that the first term of right hand side of the generic regret bound, i.e., Equation (19), is upper bounded by $O(\log(n))$. Next, noting that the second term on the right hand side of Equation (19) is analogous to the constraint violation of Algorithm 1, we observe that

$$\mathbb{E}\left[\bar{p}\left|\sum_{j=1}^{m}\left(\sum_{t=1}^{n}x_{tj}-c_j\right)\right|\right]\overset{(a)}{\leq}\mathbb{E}\left[\bar{p}\left|\sum_{j=1}^{m}\left(\sum_{t=\tau+1}^{n}x_{tj}\right)\right|\right]\overset{(b)}{\leq}\bar{p}\mathbb{E}\left[m(n-\tau)\frac{\bar{w}}{\underline{p}}\right]\overset{(c)}{\leq}\frac{m\bar{w}\bar{p}}{\underline{p}}=O(m), \tag{33}$$

where (a) follows since no constraints are violated up until the stopping time τ, (b) follows as $x_{tj}\leq\frac{\bar{w}}{\underline{p}}$, and (c) follows from Equation (25). As a result, we have established that the second term in the generic regret bound is bounded above by a constant (and thus is also bounded above by $O(\log(n))$) for Algorithm 1, which thus proves our claim.

Appendix E: Derivation of Dual of Social Optimization Problem

In this section, we derive the dual of the social optimization Problem (2a)-(2c). To this end, we first consider the following equivalent primal problem

$$\max_{\mathbf{x}_t\in\mathbb{R}^m,\,u_t}\quad U(\mathbf{x}_1,\dots,\mathbf{x}_n)=\sum_{t=1}^{n}w_t\log\left(u_t\right), \tag{34a}$$

$$\text{s.t.} \quad \sum_{t=1}^{n} x_{tj} \leq c_j, \quad \forall j \in [m], \tag{34b}$$

$$x_{tj} \geq 0, \quad \forall t \in [n], j \in [m], \tag{34c}$$

$$u_t = \sum_{j=1}^{m} u_{tj} x_{tj}, \quad \forall t \in [n], \tag{34d}$$

where we replaced the linear utility $\sum_{j=1}^{m} u_{tj} x_{tj}$ in the objective with the variable u_t and added the constraint $u_t = \sum_{j=1}^{m} u_{tj} x_{tj}$. Observe that the optimal solution of this problem is equal to that of the social optimization Problem (2a)-(2c). We now formulate the Lagrangian of this problem and derive the first order conditions of this Lagrangian to obtain the dual Problem (4).

To formulate the Lagrangian of Problem (34a)-(34d), we introduce the dual variables p_j for each good $j \in [m]$ for the capacity Constraints (34b), λ_{tj} for each user $t \in [n]$ and good $j \in [m]$ for the non-negativity Constraints (34c), and β_t for each user $t \in [n]$ for the linear utility Constraints (34d). For conciseness, we denote $\mathbf{p} \in \mathbb{R}^m$ as the vector of dual variables of the capacity Constraints (34b), $\Lambda \in \mathbb{R}^{n \times m}$ as the matrix of dual variables of the non-negativity Constraints (34c), and $\boldsymbol{\beta}$ as the vector of dual variables of the linear utility Constraints (34d). Then, we have the following Lagrangian:

$$\mathcal{L}(\{\mathbf{x}_t, u_t\}_{t=1}^{n}, \mathbf{p}, \Lambda, \boldsymbol{\beta}) = \sum_{t=1}^{n} w_t \log(u_t) - \sum_{j=1}^{m} p_j \left(\sum_{t=1}^{n} x_{tj} - c_j \right) - \sum_{t=1}^{n} \sum_{j=1}^{m} \lambda_{ij} x_{ij} - \sum_{t=1}^{n} \beta_t (u_t - \sum_{j=1}^{m} u_{tj} x_{tj})$$

Next, we observe from the first order derivative condition of the Lagrangian that

$$\frac{\partial \mathcal{L}}{\partial u_t} = \frac{w_t}{u_t} - \beta_t = 0, \quad \forall t \in [n], \text{ and}$$
$$\frac{\partial \mathcal{L}}{\partial x_{tj}} = -p_j - \lambda_{tj} + \beta_t u_{tj} = 0, \quad \forall t \in [n], j \in [m].$$

Note that we can rearrange the first equation to obtain that $u_t = \frac{w_t}{\beta_t}$ for all $t \in [n]$. Furthermore, by the sign constraint that $\lambda_{tj} \leq 0$ for all $t \in [n]$, $j \in [m]$ it follows from the second equation that $\beta_t u_{tj} \leq p_j$ for all $t \in [n]$, $j \in [m]$. Using the above equations, we can write the following dual problem:

$$\min_{\mathbf{p}, \boldsymbol{\beta}} \quad \sum_{t=1}^{n} w_t \log(w_t) - \sum_{t=1}^{n} w_t \log(\beta_t) + \sum_{j=1}^{m} p_j c_j - \sum_{t=1}^{n} w_t \tag{35}$$
$$\beta_t u_{tj} \leq p_j, \quad \forall t \in [n], j \in [m]$$

Note that at the optimal solution to the above problem $\beta_t = \min_{j \in [m]} \{ \frac{p_j}{u_{tj}} \}$. Using this observation, we can rewrite the above problem as

$$\min_{\mathbf{p}} \quad \sum_{t=1}^{n} w_t \log(w_t) - \sum_{t=1}^{n} w_t \log \left(\min_{j \in [m]} \frac{p_j}{u_{tj}} \right) + \sum_{j=1}^{m} p_j c_j - \sum_{t=1}^{n} w_t, \tag{36}$$

which is the dual Problem (4).

Appendix F: Detailed Proof Sketch of Theorem 3

The proof of Theorem 3 relies on two intermediate arguments. First, we show that if the price vector at every step of the algorithm is bounded above and below by some positive constant, then the $O(\sqrt{n})$ upper bounds on both the regret and expected constraint violation hold.

LEMMA 6 **(Regret and Constraint Violation of Algorithm 2 under Positivity and Boundedness of Prices).** *Suppose users' budget and utility parameters are drawn i.i.d. from a distribution $\mathcal{D}$ satisfying Assumption 1. Furthermore, let $\boldsymbol{\pi}$ denote the online pricing policy described by Algorithm 2, $\mathbf{x}_1, \ldots, \mathbf{x}_n$ be the corresponding allocations for the n users, and suppose that the price vector $\mathbf{p}^t$ corresponding to Algorithm 2 is such that $\mathbf{0} < \underline{\mathbf{p}} \leq \mathbf{p}^t \leq \bar{\mathbf{p}}$ for all users $t \in [n]$. Then, for a step size $\gamma = \gamma_t = \frac{\bar{D}}{\sqrt{n}}$ for some constant $\bar{D} > 0$ for all users $t \in [n]$, the regret $R_n(\boldsymbol{\pi}) \leq O(\sqrt{n})$ and the constraint violation $V_n(\boldsymbol{\pi}) \leq O(\sqrt{n})$.*

Proof Sketch. To establish this result, we proceed in three steps. First, we prove an $O(\sqrt{n})$ upper bound on the constraint violation. To do so, we sum the price update equation in Algorithm 2 across all users to establish that the excess demand for any good j is upper bounded by $\frac{p_j^{n+1}}{\gamma}$, i.e., $\sum_{t=1}^n x_{tj} - c_j \leq \frac{p_j^{n+1}}{\gamma}$. Using this relation and the fact that the prices are upper bounded by $\bar{p}$ and the step size $\gamma = O(\frac{1}{\sqrt{n}})$, we obtain the $O(\sqrt{n})$ upper bound on the constraint violation. Next, we derive a generic upper bound on the regret (different from that in the proof of Theorem 2) of any online algorithm $\boldsymbol{\pi}$ using duality (see Section 5.1), and show that $\mathbb{E}[U_n^* - U_n(\boldsymbol{\pi})] \leq \mathbb{E}\left[\sum_{t=1}^n \sum_{j=1}^m p_j^t d_j - w_t\right]$. Finally, we apply the price update rule in Algorithm 2 with a step size $\gamma = O(\frac{1}{\sqrt{n}})$ to establish an $O(\sqrt{n})$ upper bound on the term $\mathbb{E}\left[\sum_{t=1}^n \sum_{j=1}^m p_j^t d_j - w_t\right]$, i.e., the right hand side of the generic regret bound, which establishes our claim. □

We refer to Appendix G for a complete proof of Lemma 6 and note that its proof does not rely on Assumption 2.

Our second intermediary result states that if the distribution $\mathcal{D}$ satisfies Assumption 2, then the price vector $\mathbf{p}^t$ in Algorithm 2 remains strictly positive and bounded for all users $t \in [n]$.

LEMMA 7 **(Strictly Positive and Bounded Prices for Algorithm 2).** *Suppose that the budget and utility parameters of users are drawn i.i.d. from a distribution $\mathcal{D}$ satisfying Assumptions 1 and 2. Then, the price vector $\mathbf{p}^t$ corresponding to Algorithm 2 will remain strictly positive and bounded for all users $t \in [n]$ when $\gamma = \gamma_t = \frac{\bar{D}}{\sqrt{n}}$ for some constant $\bar{D} > 0$ for all users $t \in [n]$.*

Proof Sketch. To prove this claim, we proceed in two steps. First, we show that if the price vector $\mathbf{p}^t$ at each iteration of Algorithm 2 is bounded below by some vector $\underline{\mathbf{p}}$, then the price vector also remains bounded above by $\bar{\mathbf{p}}$, where $\bar{p} > 0$ is a constant, as we show in Lemma 9 in Appendix H. In other words, the positivity of prices during the operation of Algorithm 2 implies the boundedness of the prices. Next, we show that the prices of the goods will always remain positive under Assumption 2. To this end, we first consider the setting of one and two goods in the market, and then extend our analysis for the two good setting to the more general setting of m goods. We present here the main ideas to prove this result for the two good case. In particular, it directly follows from Assumption 2 that if the price of one good is small while that of another good is large, as specified by a certain price threshold p^{thresh}, then the price of the good that is small cannot become lower than $\frac{p^{\text{thresh}}\underline{u}}{2\bar{u}}$, as users will always purchase the good with the lower price given their strictly positive utilities. Next, in the case that the price of both goods is smaller than the specified threshold during the operation of Algorithm 2, we define a potential $V_t = (\mathbf{p}^t)^T \mathbf{d}$, and show that this potential is non-decreasing, i.e., $V_{t+1} \geq V_t$, if the prices of both goods are less than p^{thresh} for user t. We then use this result along with Assumption 2 to show that the price of both goods cannot go below a constant $\underline{p}$ during the operation of Algorithm 2. □

Note that Lemmas 6 and 7 jointly imply Theorem 3. For a complete proof of Lemma 7, we refer to Appendix H.

Finally, we reiterate that the key to establishing Lemma 7 lies in constructing a potential function that is non-decreasing between subsequent users when the prices of all goods are below a particular threshold during the operation of Algorithm 2. Since this result is fundamental to the proof of Lemma 7 and elucidates a close connection between the price update rule in Algorithm 2 and Fisher markets, we believe it is of independent interest. In particular, we formalize the non-decreasing potential function property of Algorithm 2 through the following lemma.

LEMMA 8 **(Non-Decreasing Potential).** *Let* $p^{thresh} = \frac{\underline{w} \min_{j \in [m]} \{d_j\}}{\sum_{j=1}^m d_j^2}$ *and define the potential* $V_t = (\mathbf{p}^t)^T \mathbf{d}$*, where* $\mathbf{p}^t$ *is the price vector corresponding to user* t *in Algorithm 2. Then, the potential for Algorithm 2 is non-decreasing for user* $t+1$*, i.e.,* $V_{t+1} \geq V_t$*.*

We refer to Appendix H.2 for a proof of Lemma 8.

Appendix G: Proof of Lemma 6

To establish this result, we proceed in three steps. First, we first prove an $O(\sqrt{n})$ upper bound on the constraint violation for the price update rule in Algorithm 2. Then, to establish an upper bound on the regret, we establish a generic bound on the regret (different from that in Lemma 1 in the proof of Theorem 2) of any online algorithm as long as the prices $\mathbf{p}_t$ are strictly positive and bounded for all users $t \in [n]$. Finally, we apply the price update rule in Algorithm 2 to establish an $O(\sqrt{n})$ upper bound on the regret for $\gamma = \gamma_t = \frac{\bar{D}}{\sqrt{n}}$ for all users $t \in [n]$ for some constant $\bar{D} > 0$.

Expected Constraint Violation Bound: To establish an $O(\sqrt{n})$ upper bound on the constraint violation, we utilize the price update rule in Algorithm 2 where $\gamma_t = \frac{\bar{D}}{\sqrt{n}}$ for some constant $\bar{D} > 0$. In particular, the price update step

$$p_j^{t+1} = p_j^t - \frac{\bar{D}}{\sqrt{n}} (d_j - x_{tj})$$

in Algorithm 2 can be rearranged to obtain

$$x_{tj} - d_j = \frac{\sqrt{n}}{\bar{D}} \left(p_j^{t+1} - p_j^t \right).$$

Summing this equation over all arriving users $t \in [n]$, it follows that

$$\sum_{t=1}^{n} x_{tj} - c_j \leq \frac{\sqrt{n}}{\bar{D}} \sum_{t=1}^{n} \left(p_j^{t+1} - p_j^t \right) = \frac{\sqrt{n}}{\bar{D}} \left(p_j^{n+1} - p_j^1 \right) \leq \frac{\sqrt{n}}{\bar{D}} p_j^{n+1} \leq \frac{\bar{p}}{\bar{D}} \sqrt{n},$$

where the last inequality follows since $p_j^{n+1} \leq \bar{p}$ by the boundedness assumption on the price vector. Using this relation, the norm of the constraint violation can be bounded as

$$\left\| \left(\sum_{t=1}^{n} \mathbf{x}_t - \mathbf{c} \right)_+ \right\|_2 \leq \left\| \sum_{t=1}^{n} \mathbf{x}_t - \mathbf{c} \right\|_2 = \sqrt{\sum_{j=1}^{m} \left(\sum_{t=1}^{n} x_{tj} - c_j \right)^2} \leq \sqrt{\sum_{j=1}^{m} \left(\frac{\bar{p}}{\bar{D}} \right)^2 n} = \sqrt{m \left(\frac{\bar{p}}{\bar{D}} \right)^2 n} \leq O(\sqrt{n}).$$

Taking an expectation of the above quantity, we obtain a $O(\sqrt{n})$ upper bound on the expected constraint violation, where $\mathbb{E}[V_n(\mathbf{x}_1, \ldots, \mathbf{x}_n)] \leq \frac{\bar{p}}{\bar{D}} \sqrt{mn} = O(\sqrt{n})$.

Generic Bound on the Regret: We now turn to establishing a generic bound on the regret of any online algorithm for which the price vector $\mathbf{p}^t$ is strictly positive and bounded for each user $t \in [n]$. To perform our analysis, let $\mathbf{p}^*$ be the optimal price vector for the following stochastic program

$$\min_{\mathbf{p}} \quad D(\mathbf{p}) = \sum_{j=1}^{m} p_j d_j + \mathbb{E} \left[\left(w \log(w) - w \log \left(\min_{j \in [m]} \frac{p_j}{u_j} \right) - w \right) \right]. \tag{37}$$

Then, by duality we have that the primal objective value U_n^* is no more than the dual objective value with $\mathbf{p} = \mathbf{p}^*$, which gives the following upper bound on the optimal objective

$$U_n^* \leq \sum_{j=1}^{m} p_j^* c_j + \sum_{t=1}^{n} \left(w_t \log(w_t) - w_t \log \left(\min_{j \in [m]} \frac{p_j^*}{u_{tj}} \right) - w_t \right).$$

Then, taking an expectation on both sides of the above inequality, it follows that

$$\begin{aligned} \mathbb{E}[U_n^*] &\leq \mathbb{E} \left[\sum_{j=1}^{m} p_j^* c_j + \sum_{t=1}^{n} \left(w_t \log(w_t) - w_t \log \left(\min_{j \in [m]} \frac{p_j^*}{u_{tj}} \right) - w_t \right) \right], \\ &= nD(\mathbf{p}^*), \end{aligned}$$

by the definition of $D(\mathbf{p})$ in Problem (37). Finally, noting that $\mathbf{p}^*$ is the optimal solution to the stochastic Program (37), it follows that

$$\begin{aligned} \mathbb{E}\left[U_n^*\right] \leq nD(\mathbf{p}^*) &\overset{(a)}{\leq} \sum_{t=1}^{n} \mathbb{E}\left[D(\mathbf{p}^t)\right] \overset{(b)}{=} \sum_{t=1}^{n} \mathbb{E} \left[\sum_{j=1}^{m} p_j^t d_j + w_t \log(w_t) - w_t \log \left(\min_{j \in [m]} \frac{p_j^t}{u_{tj}} \right) - w_t \right], \\ &\overset{(c)}{=} \mathbb{E} \left[\sum_{t=1}^{n} \left(\sum_{j=1}^{m} p_j^t d_j + w_t \log(w_t) - w_t \log \left(\min_{j \in [m]} \frac{p_j^t}{u_{tj}} \right) - w_t \right) \right], \end{aligned}$$

where (a) follows by the optimality of $\mathbf{p}^*$ for the stochastic Program (37), (b) follows by the definition of $D(\mathbf{p}^t)$, and (c) follows from the linearity of expectations.

Next, let j_t be a good in the optimal consumption set S_t^* for user t given the price vector $\mathbf{p}^t$. Then, the true accumulated social welfare objective under an algorithm $\boldsymbol{\pi}$ can be expressed as

$$\begin{aligned} U_n(\boldsymbol{\pi}) &= \sum_{t=1}^n w_t \log\left(\sum_{j=1}^m u_{tj} x_{tj}\right), \\ &= \sum_{t=1}^n w_t \log\left(\sum_{j=1}^m u_{tj} \mathbb{1}_{j=j_t} \frac{w_t}{p_j^t}\right), \end{aligned}$$

which follows since the utility when consuming any feasible bundle of goods in their optimal consumption set equals their utility when purchasing $\frac{w_t}{p_{j_t}^t}$ units of good $j_t \in S_t^*(\mathbf{p}^t)$. Finally combining the upper bound on the expected optimal objective and the above obtained relation on the accumulated objective under an algorithm $\boldsymbol{\pi}$, we obtain the following upper bound on the expected regret

$$\mathbb{E}[U_n^* - U_n(\boldsymbol{\pi})] \leq \mathbb{E}\left[\sum_{t=1}^n \left(\sum_{j=1}^m p_j^t d_j + w_t \log(w_t) - w_t \log\left(\min_{j \in [m]} \frac{p_j^t}{u_{tj}}\right) - w_t\right)\right] \tag{38}$$

$$- \mathbb{E}\left[\sum_{t=1}^n w_t \log\left(u_{tj_t} \frac{w_t}{p_{j_t}^t}\right)\right], \tag{39}$$

$$= \mathbb{E}\left[\sum_{t=1}^n \sum_{j=1}^m p_j^t d_j - w_t\right], \tag{40}$$

where the final equality follows as $j_t \in S_t^*(\mathbf{p}^t)$.

Square Root Regret Bound: We now use the generic regret bound derived in Equation (40) for any online algorithm with bounded prices that are always strictly positive for each $t \in [n]$ to obtain an $O(\sqrt{n})$ upper bound on the regret of Algorithm 2. In particular, we use the price update equation in Algorithm 2 to derive the $O(\sqrt{n})$ regret bound. We begin by observing from the price update equation that

$$\left\|\mathbf{p}^{t+1}\right\|_2^2 = \left\|\mathbf{p}^t - \frac{\bar{D}}{\sqrt{n}}\left(\mathbf{d} - \mathbf{x}_t\right)\right\|_2^2.$$

Expanding the right hand side of the above equation, we obtain that

$$\left\|\mathbf{p}^{t+1}\right\|^2 \leq \left\|\mathbf{p}^t\right\|^2 - \frac{2\bar{D}}{\sqrt{n}}\left(\sum_{j=1}^m p_j^t d_j - \sum_{j=1}^m p_j^t x_{tj}\right) + \frac{\bar{D}^2}{n}\left\|\mathbf{d} - \mathbf{x}_t\right\|^2.$$

We can then rearrange the above equation to obtain

$$\sum_{j=1}^m p_j^t d_j - \sum_{j=1}^m p_j^t x_{tj} \leq \frac{\sqrt{n}}{2\bar{D}}\left(\left\|\mathbf{p}^t\right\|^2 - \left\|\mathbf{p}^{t+1}\right\|^2\right) + \frac{\bar{D}}{2\sqrt{n}}\left\|\mathbf{d} - \mathbf{x}_t\right\|^2$$

Finally, summing both sides of the above equation over $t \in [n]$, we get

$$\sum_{t=1}^{n}\sum_{j=1}^{m} p_j^t d_j - \sum_{t=1}^{n}\sum_{j=1}^{m} p_j^t x_{tj} \leq \frac{\sqrt{n}}{2\bar{D}} \sum_{t=1}^{n} \left(\left\| \mathbf{p}^t \right\|^2 - \left\| \mathbf{p}^{t+1} \right\|^2 \right) + \sum_{t=1}^{n} \frac{\bar{D}}{2\sqrt{n}} \left\| \mathbf{d} - \mathbf{x}_t \right\|_2^2, \tag{41}$$

$$\stackrel{(a)}{\leq} \frac{\sqrt{n}}{2\bar{D}} \left\| p^1 \right\|^2 + \frac{\bar{D}}{2\sqrt{n}} \sum_{t=1}^{n} m \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}} \right)^2, \tag{42}$$

$$\leq \sqrt{n} \left(\frac{\left\| p^1 \right\|^2}{2\bar{D}} + \frac{\bar{D} m}{2} \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}} \right)^2 \right), \tag{43}$$

$$\leq O(\sqrt{n}), \tag{44}$$

where the (a) follows by the boundedness of the consumption vector for each agent, since the prices are strictly positive and bounded below by $\underline{p} > 0$. Finally, noting that all agents completely spend their budget at the optimal solution of the individual optimization problem, i.e., $\sum_{j \in [M]} p_j^t x_{tj} = w_t$, we obtain from the generic regret bound in Equation (40) that

$$\mathbb{E}[U_n^* - U_n(\boldsymbol{\pi})] \leq \mathbb{E}\left[\sum_{t=1}^{n}\sum_{j=1}^{m} p_j^t d_j - w_t \right] = \sum_{t=1}^{n} \mathbb{E}\left[\sum_{j=1}^{m} p_j^t d_j - \sum_{t=1}^{n}\sum_{j=1}^{m} p_j^t x_{tj} \right],$$

$$\stackrel{(a)}{\leq} \sqrt{n} \left(\frac{\left\| p^1 \right\|^2}{2\bar{D}} + \frac{\bar{D} m}{2} \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}} \right)^2 \right),$$

$$= O(\sqrt{n}),$$

where (a) follows from Equation (43). Thus, we have proven the $O(\sqrt{n})$ upper bound on the expected regret of Algorithm 2 under the assumed conditions on the price vectors $\mathbf{p}^t$ for all users $t \in [n]$.

Appendix H: Proof of Lemma 7

To establish this result, we proceed in two steps. In particular, we first show that the strict positivity of prices during the operation of Algorithm 2 implies that the prices are bounded for all $t \in [n]$ in Appendix H.1. Then, we show that the prices of the goods will always remain positive under Assumption 2 in Appendix H.2.

H.1. Positivity of Prices Implies Boundedness

We show through the following lemma that if the price vector $\mathbf{p}^t$ is bounded below by some vector $\underline{\mathbf{p}}$ at each iteration of Algorithm 2, then the price vector also remains bounded above by $\bar{\mathbf{p}}$, where each component $\bar{p} > 0$ of the vector $\bar{\mathbf{p}}$ is a constant.

LEMMA 9 **(Positivity Implies Price Boundedness in Algorithm 2).** *Suppose that the budget and utility parameters of users are drawn i.i.d. from a distribution $\mathcal{D}$ satisfying Assumption 1, and the price vector $\mathbf{p}^t \geq \underline{\mathbf{p}} > \mathbf{0}$ for all users $t \in [n]$. Then, the price vector $\mathbf{p}^t$ corresponding to Algorithm 2 is bounded at each time an agent $t \in [n]$ arrives, i.e., $\mathbf{p}^t \leq \bar{\mathbf{p}}$ for all $t \in [n]$ for some vector $\bar{\mathbf{p}} \geq \underline{\mathbf{p}}$, when the step-size $\gamma = \gamma_t = \frac{\bar{D}}{\sqrt{n}}$ for some $0 < \bar{D} \leq 1$.*

Proof. We establish that the prices of all goods are always bounded above at each step of Algorithm 2 if the prices of the goods are bounded below by $\underline{\mathbf{p}}$ at each step. To show that the prices are bounded above, we consider the settings when (i) $\|\mathbf{p}^t\|_2 \geq \frac{m\left(\bar{d}+\frac{\bar{w}}{\underline{p}}\right)^2+2\bar{w}}{2\underline{d}}$, and (ii) $\|\mathbf{p}^t\|_2 \leq \frac{m\left(\bar{d}+\frac{\bar{w}}{\underline{p}}\right)^2+2\bar{w}}{2\underline{d}}$, where $\bar{d} = \max_{j \in [m]} d_j$ and $\underline{d} = \min_{j \in [m]} d_j > 0$. In case (i), we observe that

$$
\begin{aligned}
\left\|\mathbf{p}^{t+1}\right\|_2^2 = \left\|\mathbf{p}^t - \gamma(\mathbf{d}-\mathbf{x}_t)\right\|_2^2 &= \left\|\mathbf{p}^t\right\|_2^2 - 2\gamma(\mathbf{p}^t)^\top(\mathbf{d}-\mathbf{x}^t) + \gamma^2\left\|\mathbf{d}-\mathbf{x}^t\right\|_2^2, \\
&\overset{(a)}{\leq} \left\|\mathbf{p}^t\right\|_2^2 + 2\gamma\bar{w} + \gamma^2 m\left(\bar{d}+\frac{\bar{w}}{\underline{p}}\right)^2 - 2\gamma(\mathbf{p}^t)^\top\mathbf{d}, \\
&\leq \left\|\mathbf{p}^t\right\|_2^2 + 2\gamma\bar{w} + \gamma^2 m\left(\bar{d}+\frac{\bar{w}}{\underline{p}}\right)^2 - 2\gamma\underline{d}\left\|\mathbf{p}^t\right\|_1, \\
&\overset{(b)}{\leq} \left\|\mathbf{p}^t\right\|_2^2 + 2\gamma\bar{w} + \gamma^2 m\left(\bar{d}+\frac{\bar{w}}{\underline{p}}\right)^2 - 2\gamma\underline{d}\left\|\mathbf{p}^t\right\|_2, \\
&\overset{(c)}{\leq} \left\|\mathbf{p}^t\right\|_2^2,
\end{aligned}
$$

where (a) follows from the fact that $(\mathbf{p}^t)^\top\mathbf{x}^t = w_t \leq \bar{w}$ and $\mathbf{x}^t \leq \frac{\bar{w}}{\underline{p}}\mathbf{1}$, where $\mathbf{1}$ is an m-dimensional vector of all ones, (b) follows from the norm equivalence relation between the one and the two norms, and (c) follows for any step-size $\gamma \leq 1$ and the fact that $\|\mathbf{p}^t\|_2 \geq \frac{m\left(\bar{d}+\frac{\bar{w}}{\underline{p}}\right)^2+2\bar{w}}{2\underline{d}}$.

Next, in case (ii) it holds that

$$
\begin{aligned}
\left\|\mathbf{p}^{t+1}\right\|_2 = \left\|\mathbf{p}^t - \gamma(\mathbf{d}-\mathbf{x}_t)\right\|_2 &\overset{(a)}{\leq} \left\|\mathbf{p}^t\right\|_2 + \gamma\left\|\mathbf{d}-\mathbf{x}_t\right\|_2 \overset{(b)}{\leq} \left\|\mathbf{p}^t\right\|_2 + \gamma\left\|\mathbf{d}-\mathbf{x}_t\right\|_1, \\
&\overset{(c)}{\leq} \frac{m\left(\bar{d}+\frac{\bar{w}}{\underline{p}}\right)^2+2\bar{w}}{2\underline{d}} + m\left(\bar{d}+\frac{\bar{w}}{\underline{p}}\right),
\end{aligned}
$$

where (a) follows by the triangle inequality, (b) follows from the norm equivalence relation between the one and to norms, and (c) holds since $\|\mathbf{p}^t\|_2 \leq \frac{m\left(\bar{d}+\frac{\bar{w}}{\underline{p}}\right)^2+2\bar{w}}{2\underline{d}}$ and $\mathbf{x}^t \leq \frac{\bar{w}}{\underline{p}}\mathbf{1}$.

From the above inequalities, we observe that $\|\mathbf{p}^t\|_2 \leq \frac{m\left(\bar{d}+\frac{\bar{w}}{\underline{p}}\right)^2+2\bar{w}}{2\underline{d}} + m(\bar{d}+\frac{\bar{w}}{\underline{p}})$ for all t. This relation implies that the price vector of Algorithm 2 is always bounded above when the price vector of Algorithm 2 is bounded below by $\underline{\mathbf{p}}$ at each step, which completes the proof of Lemma 9.

H.2. Positivity of Prices

We now show that under Assumption 2 the prices of the goods remain strictly positive during the operation of Algorithm 2. To this end, we first prove this claim for the setting of one good. Then, we leverage Assumption 2 to construct an argument for the setting of two goods and finally, extend this argument to the general setting of m goods.

H.2.1. One Good Suppose that there is exactly one good in the market with a capacity c_1. Then, we claim that under Algorithm 2, the price of each good for each user $t \in [n]$ remains bounded between $\underline{p}$ and $\bar{p}$, where

$\underline{p} > 0$ and $\bar{p}$ are constants. In particular, we prove this claim for $\underline{p} = \frac{\underline{w}}{2d_1}$, and $\bar{p} = \frac{\bar{w}}{d_1} + d_1 + \frac{\bar{w}d_1}{\underline{w}}$ that $\underline{p} \leq p_1^t \leq \bar{p}$ for all $t \in [n]$, when $\gamma = \frac{\bar{D}}{\sqrt{n}}$, where $\bar{D} = \frac{\frac{\underline{w}}{2d_1}}{d_1 + \frac{\bar{w}d_1}{\underline{w}}}$, and the initial price vector $p_1^1 \in [\frac{\underline{w}}{2d_1}, \frac{\bar{w}}{d_1} + d_1 + \frac{\bar{w}d_1}{\underline{w}}]$.

To this end, first observe by the price update rule that

$$p_1^{t+1} = p_1^t - \gamma\left(d_1 - x_{t1}\right) = p_1^t - \gamma\left(d_1 - \frac{w_t}{p_1^t}\right)$$

since it is optimal for each user to purchase $\frac{w_t}{p_1^t}$ units of good one as there is only one good in the market. Next, to establish the the bounds on the price, we show that if p_1^t is large (small), then it must hold that $p_1^{t+1} \leq p_1^t$ ($p_1^{t+1} \geq p_1^t$). In particular, observe that if $p_1^t \geq \frac{\bar{w}}{d_1}$, then

$$p_1^{t+1} = p_1^t - \gamma(d_1 - \frac{w_t}{p_1^t}) \leq p_1^t - \gamma d_1 + \gamma\frac{w_t d_1}{\bar{w}} \leq p_1^t.$$

and if $p_1^{t+1} \leq \frac{\underline{w}}{d_1}$, then $p_1^{t+1} \geq p_1^t$. On the other hand, when $\frac{\underline{w}}{d_1} \leq p_1^t \leq \frac{\bar{w}}{d_1}$, then it holds that

$$p_1^{t+1} = p_1^t - \gamma(d_1 - \frac{w_t}{p_1^t}) \leq \frac{\bar{w}}{d_1} + d_1 + \frac{\bar{w}d_1}{\underline{w}},$$

since $\gamma \leq 1$, which holds true for large n. This establishes the upper bound on the price. For the lower bound on the price, observe that when $\frac{\underline{w}}{d_1} \leq p_1^t \leq \frac{\bar{w}}{d_1}$, then it holds that

$$p_1^{t+1} = p_1^t - \gamma(d_1 - \frac{w_t}{p_1^t}) \geq \frac{\underline{w}}{d_1} - \gamma(d_1 + \frac{\bar{w}d_1}{\underline{w}}) \geq \frac{\underline{w}}{2d_1},$$

since $\gamma = \frac{\bar{D}}{\sqrt{n}}$, where $\bar{D} = \frac{\frac{\underline{w}}{2d_1}}{d_1 + \frac{\bar{w}d_1}{\underline{w}}}$. This establishes the lower bound on the price, which proves our claim that for the setting of one good the price for each user $t \in [n]$ under Algorithm 2 is always bounded away from zero and is bounded above by some constant.

H.2.2. Two Goods We now consider the setting where there are two goods in the market with capacities c_1 and c_2, respectively. Furthermore, we consider the setting when the support of the utilities and budgets of users is strictly positive. In particular, let $\rho = \frac{\bar{u}}{\underline{u}}$ be the maximum ratio of the utilities in the support of the distribution $\mathcal{D}$. In this case, we claim that under Algorithm 2, the price of both goods for each user $t \in [n]$ is always bounded away from zero. To establish this claim, we analyse several cases for when the prices of the goods is above or below certain thresholds. In each of these cases, we show that the prices of both goods for each user are always bounded away from zero by some constant $\tilde{\underline{p}} = \frac{1}{(4\rho)^2}\frac{\underline{w}\min\{d_1,d_2\}}{d_1^2+d_2^2}$.

To prove our claim, we first let the initial price vector be $p_j^1 \in \left[\frac{\underline{w}d_j}{\sum_{j=1}^2 d_j^2}, \frac{\bar{w}d_j}{\sum_{j=1}^2 d_j^2}\right]$. Next, observe that at each time when a user t arrives, the prices of the goods in the market must fall within one of the following cases:

1. Case 1 (Both Prices are Large): $p_1^t \geq \frac{\bar{w}d_1}{d_1^2+d_2^2}$ and $p_2^t \geq \frac{\bar{w}d_2}{d_1^2+d_2^2}$;
2. Case 2 (Both Prices are Small): $p_1^t < \frac{\underline{w}d_1}{d_1^2+d_2^2}$ and $p_2^t < \frac{\underline{w}d_2}{d_1^2+d_2^2}$;
3. Case 3 (Intermediate Prices): One of the prices $p_j^t < \frac{\underline{w}d_j}{d_1^2+d_2^2}$ while the other price $p_{j'}^t \geq \frac{\underline{w}d_{j'}}{d_1^2+d_2^2}$

Here case 1 corresponds to the setting when the prices of both goods is large, while case 2 represents the setting when the prices of both goods is small. On the other hand, case 3 captures all the other intermediate cases, where either both prices are bounded above and below or one of the prices is bounded above and the other is bounded below. To show that the prices are bounded below by the above defined $\underline{\tilde{p}}$, we first note that in Case (i) that the prices of both goods are large and since the amount by which prices can drop is at most $O(\frac{1}{\sqrt{T}})$ at each step that the price at step $t+1$ will clearly be bounded below by $\underline{\tilde{p}}$. Next, observe in case 3 that since the price of one of the goods is above the specified threshold. Without loss of generality, suppose that this is good one. Then, by the boundedness of utilities, we know that the price of good two must be at least $\frac{1}{2\rho}\frac{\underline{w}d_1}{d_1^2+d_2^2}$, as users would only purchase good two if their price fell below $\frac{1}{\rho}\frac{\underline{w}d_1}{d_1^2+d_2^2}$. Thus, if we are in case 3, then it must also be that the prices of both goods are bounded below by $\underline{\tilde{p}}$. Finally, we proceed to analysing case 2.

Analysis of Case 2: In case 2, we establish that the prices of the goods at each time a user arrives is bounded below in the setting when the good prices are low. In particular, we now show that both the good prices at step $t+1$ can be no lower than $\underline{\tilde{p}}$. To this end, we proceed in the following steps. First, we define a potential $V_t = (\mathbf{p}^t)^\top \mathbf{d}$, and show that this potential is non-decreasing, i.e., $V_{t+1} \geq V_t$ if the price vector $\mathbf{p}^t$ satisfies the condition of case two. This claim establishes that the potential V_t is monotonically non-decreasing until the price vector $\mathbf{p}^t$ exits case two. Then, we use this result and the fact that the utilities and budgets of users are bounded below to establish that the prices of the two goods are always strictly positive.

Proof of $V_{t+1} \geq V_t$ in Case 2: Let $V_t = p_1^t d_1 + p_2^t d_2$, where $p_1^t < \frac{\underline{w}d_1}{d_1^2+d_2^2}$ and $p_2^t < \frac{\underline{w}d_2}{d_1^2+d_2^2}$. Then, we show that $V_{t+1} \geq V_t$. To see this, suppose, without loss of generality, that buyer t consumes $\frac{aw_t}{p_1^t}$ units of good one and $\frac{(1-a)w_t}{p_2^t}$ units of good two, where $a \in [0,1]$. Then the price of good one for user $t+1$ is given by

$$p_1^{t+1} = p_1^t - \gamma\left(d_1 - \frac{aw_t}{p_1^t}\right) \geq p_1^t - \gamma\left(d_1 - \frac{aw_t(d_1^2+d_2^2)}{\underline{w}d_1}\right) \geq p_1^t - \gamma\left(d_1 - \frac{a(d_1^2+d_2^2)}{d_1}\right).$$

Similarly, the price of good two for user $t+1$ is given by

$$p_2^{t+1} = p_2^t - \gamma\left(d_2 - \frac{(1-a)w_t}{p_2^t}\right) \geq p_2^t - \gamma\left(d_2 - \frac{(1-a)(d_1^2+d_2^2)}{d_2}\right).$$

Using the above inequalities for the prices of the two goods for user $t+1$, we obtain that $V_{t+1} \geq V_t$ since

$$\begin{aligned}
V_{t+1} &= p_1^{t+1}d_1 + p_2^{t+1}d_2, \\
&\geq p_1^t d_1 + p_2^t d_2 - \gamma d_1\left(d_1 - \frac{a(d_1^2+d_2^2)}{d_1}\right) - \gamma d_2\left(d_2 - \frac{(1-a)(d_1^2+d_2^2)}{d_2}\right), \\
&= V_t - \gamma\left(d_1^2 - a(d_1^2+d_2^2) + d_2^2 - (1-a)(d_1^2+d_2^2)\right), \\
&= V_t,
\end{aligned}$$

which proves our claim that the potential is non-decreasing when the price vector $\mathbf{p}^t$ lies in case two.

V_t forms a monotonic sequence in case 2: We have observed that the potential is non-decreasing for each user t when $\mathbf{p}^t$ is in case two. Now, let τ_1 be the index of the first user when the price vector belongs to case two, and it holds that $\tau_2 > \tau_1$ is the user index for which the price vector exits case 2 or at which the algorithm ends, i.e., $\tau_2 = \min\{\{t > \tau_1 : p_1^t \geq \frac{\underline{w} d_1}{d_1^2 + d_2^2} \text{ or } p_2^t \geq \frac{\underline{w} d_2}{d_1^2+d_2^2}\}, n+1\}$. Then, from the above analysis that $V_{t+1} \geq V_t$ if the price $\mathbf{p}^t$ is in case two, it holds that

$$V_{\tau_1} \leq V_{\tau_1+1} \leq \ldots \leq V_{\tau_2}.$$

Prices of both goods are strictly positive in Case 2: Since τ_1 is the first user index for which the prices of the goods belongs to case two, it must hold that the price of at least one of the goods exceeds the respective threshold for user $\tau_1 - 1$. Without loss of generality, suppose that $p_1^{\tau_1 - 1} \geq \frac{\underline{w} d_1}{d_1^2+d_2^2}$.

Next, observe that the price of good one for user τ_1 must be such that

$$p_1^{\tau_1} = p_1^{\tau_1 - 1} - \gamma(d_1 - x_{\tau_1 - 1, 1}) \geq \frac{\underline{w} d_1}{d_1^2 + d_2^2} - \gamma d_1.$$

Since we can take $\gamma \leq \frac{\frac{\underline{w} d_1}{d_1^2+d_2^2}}{2d_1}$, it follows that $p_1^{\tau_1} \geq \frac{\underline{w} d_1}{2(d_1^2+d_2^2)}$.

Furthermore, since the utilities are bounded below, it follows that $p_2^{\tau_1} \geq \frac{1}{2\rho} p_1^{\tau_1} \geq \frac{1}{4\rho} \frac{\underline{w} d_1}{2(d_1^2+d_2^2)}$, as good two will be the only one consumed when its price is lower than $\frac{1}{\rho}$ times the price of good one. We now show that at all points between τ_1 and $\tau_2 - 1$ that the prices of both goods is bounded below by $\underline{\tilde{p}}$. To see this, first note that the price of good two must always be at least $\frac{1}{4\rho} \frac{\underline{w} d_1}{2(d_1^2+d_2^2)} \geq \underline{\tilde{p}}$ by the monotonicity property of the potential function (as at least one of the good prices must increase but the price of good two cannot fall below $\frac{1}{4\rho} \frac{\underline{w} d_1}{2(d_1^2+d_2^2)}$ by the boundedness of utilities below). Analogously, since the price of good two cannot fall below $\frac{1}{4\rho} \frac{\underline{w} d_1}{2(d_1^2+d_2^2)}$ between τ_1 and τ_2, it also follows that the price of good one cannot fall below $\frac{1}{8\rho^2} \frac{\underline{w} \min\{d_1, d_2\}}{2(d_1^2+d_2^2)} \geq \underline{\tilde{p}}$ by the the boundedness of utilities below. Thus, we have established that both $p_1^t, p_2^t \geq \underline{\tilde{p}}$ for all users $t \in \{\tau_1, \ldots \tau_2\}$. We note that we can repeat the above line of reasoning for all periods when the price vector belongs to case 2 and thus have shown that the prices are bounded below by a constant, which establishes our claim.

H.2.3. Extending Above Argument to Multiple Goods To extend our analysis from the setting of two goods to that of m goods, we first let $\mathbf{p}_j^1 \in \left[\frac{\underline{w} d_j}{\sum_{j \in [m]} d_j^2}, \frac{\bar{w} d_j}{\sum_{j \in [m]} d_j^2}\right]$ for all goods j. Then, we also consider multiple cases as in the two-good setting, and observe that to establish a lower bound on the prices, by the boundedness of utilities from below it suffices to consider the case when all goods have a price strictly below $\frac{\underline{w} d_j}{\sum_{j \in [m]} d_j^2}$.

To show that the prices are also bounded by below, we follow a similar line of reasoning as for the two good case. To this end, first observe that as in the two good case, there must be a good j (without loss of generality suppose this is good one) that has a price just below $\frac{\underline{w} d_1}{\sum_{j \in [m]} d_j^2}$, e.g., $\frac{\underline{w} d_1}{2\sum_{j \in [m]} d_j^2}$, when the price vector satisfies the condition that all goods have a price strictly below $\frac{\underline{w} d_1}{\sum_{j \in [m]} d_j^2}$ for the first time.

This implies by by the boundedness of utilities from below that the price of the other goods are at least $p_2 \geq \frac{\underline{w} d_1}{4\rho \sum_{j \in [m]} d_j^2}$, $p_3 \geq \frac{\underline{w} d_1}{(4\rho)^2 \sum_{j \in [m]} d_j^2}$, ..., $p_m \geq \frac{\underline{w} d_j}{(4\rho)^{m-1} \sum_{j \in [m]} d_j^2}$.

Next, it can again be shown in the multiple good case that the potential $V_{t+1} = (\mathbf{p}^{t+1})^T \mathbf{d} \geq (\mathbf{p}^t)^T \mathbf{d} = V_t$ in the case when all goods j have a price strictly below $\frac{\underline{w} d_j}{\sum_{j \in [m]} d_j^2}$ at each time a user t arrives. Using the fact that $V_{t+1} \geq V_t$, it follows that the price of one of the goods must always be above their respective lower bounds. However, the price of good m must be above its threshold since otherwise we would violate the boundedness of utilities from below. Since the price of good m is bounded from below by a positive constant it follows that the prices of all the other goods must be at least $\frac{\underline{w} d_j}{(4\rho)^{2m} \sum_{j \in [m]} d_j^2}$ by the boundedness of utilities from below. This completes the claim that the prices of the goods are always bounded below, which proves our claim.

Appendix I: Remarks on the Positivity of Prices in Algorithm 2

In this section, we show that the price vector $\mathbf{p}^t$ is strictly positive during the operation of Algorithm 2 for all $t \in [n]$ with high probability for distributions $\mathcal{D}$ such as in the counterexample used to prove Theorem 1.

In particular, we consider the class of distributions $\mathcal{D}$ that satisfy the following natural assumption, which states that the expected consumption of a good by any user is strictly greater than their market share of that good if the price of the good is small.

ASSUMPTION 3. *There exists $\underline{\tilde{p}}$ such that if $p_j < \underline{\tilde{p}}$ for any good j, then the distribution $\mathcal{D}$ is such that the expected consumption of that good is at least $\frac{d_j}{1-\delta}$ for some $\delta > 0$.*

We note that Assumption 3 imposes a mild restriction on the set of allowable distributions from which the utility parameters of users are drawn. In particular, the assumption on the distribution $\mathcal{D}$ implies that for each good there are a certain fraction of the arriving users with a sufficiently high utility for that good. As a result if the price of a good drops too low then a certain fraction of users will purchase large quantities of that good that is far greater than their market share for that good. For instance, the distribution $\mathcal{D}$ constructed in the counterexample in the proof of Theorem 1 satisfies Assumption 3, as the expected consumption of each good is strictly greater than each user's market share d_j of that good if its price drops strictly below 0.5. As a result, Assumption 3 intuitively implies that the price of any good cannot drop "too far" below some specified price $\underline{\tilde{p}}$ during the operation of Algorithm 2.

We now apply the Chernoff bound and use Assumption 3 to claim that the price vector $\mathbf{p}^t$ for all users $t \in [n]$ is lower bounded by $\underline{\mathbf{p}}$ with high probability for some constant $\underline{p} > 0$. To this end, as in Assumption 3, let $\underline{\tilde{p}}$ be a constant and let $0 < \underline{p} \leq (1-\epsilon)(\underline{\tilde{p}} - \gamma d_j)$ for small $\epsilon > 0$. We now provide a bound on the probability that the price p_j^t of some good j for some user t drops below $\underline{p}$ during the operation of Algorithm 2. Here, we assume that the initial price $\mathbf{p}^1$ in Algorithm 2 is sufficiently higher than $\underline{\mathbf{p}}$. We now suppose that t is the first time step at which the price of some good j falls below $\underline{\tilde{p}}$, and that the price of that good stays below

$\underline{\tilde{p}}$ for another k steps. Then, we can upper bound the probability that the price p_j^t of some good j for some user t drops below $\underline{p}$ as follows

$$\begin{aligned}
\mathbb{P}[p_j^t \leq \underline{p} \text{ for some } j, t] &\overset{(a)}{\leq} \mathbb{P}[p_j^t \leq (1-\epsilon)(\underline{\tilde{p}} - \gamma d_j) \text{ for some good } j \text{ for some user } t], \\
&\overset{(b)}{\leq} \mathbb{P}\left[\underline{\tilde{p}} - \gamma d_j - \gamma k d_j + \gamma \sum_{t'=t}^{t+k-1} x_{t'j} \leq (1-\epsilon)(\underline{\tilde{p}} - \gamma d_j) \text{ for some } j, t\right], \\
&\overset{(c)}{=} \mathbb{P}\left[\gamma \sum_{t'=t}^{t+k-1} x_{t'j} \leq -\epsilon(\underline{\tilde{p}} - \gamma d_j) + \gamma k d_j \text{ for some } j, t\right], \\
&\overset{(d)}{\leq} \mathbb{P}\left[\sum_{t'=t}^{t+k-1} x_{t'j} \leq k d_j \text{ for some } j, t\right],
\end{aligned}$$

where (a) follows as $\underline{p} \leq (1-\epsilon)(\underline{\tilde{p}} - \gamma d_j)$ for small $\epsilon > 0$, (b) follows by the price update rule in Algorithm 2 and the fact that t is the first step at which the price of some good j falls below $(\underline{\tilde{p}} - \gamma d_j)$, (c) follows by rearranging the terms in the inequality, and (d) follows as $-\epsilon(\underline{\tilde{p}} - \gamma d_j) < 0$. To upper bound the right hand side term $\mathbb{P}[\sum_{t'=t-1}^{t+k-1} x_{t'j} \leq (k+1)d_j \text{ for some } j, t]$, we begin by noting that the user consumption $x_{t'j}$ is not i.i.d. since user's consumption bundles depend on the price, which is inherently dependent on the budget and utility parameters of earlier users by the price update equation of Algorithm 2. However, defining $\mathbb{E}_{t-1}[\cdot] = \mathbb{E}[\cdot|(w_1, \mathbf{u}_1), \ldots, (w_1, \mathbf{u}_{t-1})]$ as the conditional expectation of the allocations of Algorithm 2 depending on the realizations of the users' parameters, we can rewrite $\mathbb{P}[\sum_{t'=t-1}^{t+k-1} x_{t'j} \leq (k+1)d_j \text{ for some } j, t]$ as follows:

$$\begin{aligned}
\mathbb{P}\left[\sum_{t'=t}^{t+k-1} x_{t'j} \leq k d_j \text{ for some } j, t\right] &= \mathbb{P}\left[\sum_{t'=t}^{t+k-1} x_{t'j} - \sum_{t'=t}^{t+k-1} \mathbb{E}_{t'-1}[x_{t'j}] \leq k d_j - \sum_{t'=t}^{t+k-1} \mathbb{E}_{t'-1}[x_{t'j}] \text{ for some } j, t\right], \\
&\leq \mathbb{P}\left[\sum_{t'=t}^{t+k-1} x_{t'j} - \sum_{t'=t}^{t+k-1} \mathbb{E}_{t'-1}[x_{t'j}] \leq k d_j - k\frac{d_j}{1-\delta} \text{ for some } j, t\right], \\
&= \mathbb{P}\left[\sum_{t'=t}^{t+k-1} x_{t'j} - \sum_{t'=t}^{t+k-1} \mathbb{E}_{t'-1}[x_{t'j}] \leq -\frac{k d_j \delta}{1-\delta} \text{ for some } j, t\right],
\end{aligned}$$

where the inequality follows by Assumption 3 that $\mathbb{E}_{t'-1}[x_{t'j}] \geq \frac{d_j}{1-\delta}$ for all $t \leq t' \leq t+k-1$ as $p_j^{t'} \leq \underline{\tilde{p}}$ for this range of values of t'. Furthermore, defining $\kappa_{t'j} = x_{t'j} - \mathbb{E}_{t'-1}[x_{t'j}]$ and $S_n = \sum_{t'=t}^{t+n-1} \kappa_{t'j}$, we note that S_n is a martingale with respect to the filtration $\sigma((w_1, \mathbf{u}_1), \ldots, (w_t, \mathbf{u}_t), \ldots, (w_{t+n-1}, \mathbf{u}_{t+n-1}))$. To see this, observe that

$$\mathbb{E}[S_{n+1}|S_1, \ldots, S_n] = S_n + \mathbb{E}[\kappa_{t+n-1,j}] = S_n + \mathbb{E}[x_{t+n-1,j} - \mathbb{E}_{t+n-2}[x_{t+n-1,j}]] = S_n.$$

noting that k is the number of steps for which the price of a good j remains below $\underline{\tilde{p}}$ and that the step size is $O(\frac{1}{\sqrt{n}})$, it follows that for a constant reduction in the price of good j, i.e., for $p_j^{t'} \leq (1-\epsilon)\underline{\tilde{p}}$, it must hold that $k = O(\sqrt{n})$. Thus, for any $k = o(\sqrt{n})$, it must hold that $x_{t'j}$ is bounded for all $t \leq t' \leq t+k-1$

(as the price remains strictly positive for $k = o(\sqrt{n})$) and thus the corresponding martingale has bounded differences, i.e., $|S_n - S_{n-1}| \leq L$ for some constant L. Then, by Azuma's inequality for martingales with bounded differences (Lalley 2013), it follows that

$$\mathbb{P}\left[\sum_{t'=t}^{t+k-1} x_{t'j} - \sum_{t'=t}^{t+k-1} \mathbb{E}_{t'-1}[x_{t'j}] \leq -\frac{kd_j\delta}{1-\delta} \text{ for some } j, t\right] \leq e^{\frac{-k^2 d_j^2 \delta^2}{2kL^2(1-\delta)^2}} = e^{\frac{-kd_j^2\delta^2}{2L^2(1-\delta)^2}}$$

for $k = o(\sqrt{n})$. Finally, combining the above derived sequence of inequalities, we obtain that

$$\mathbb{P}[p_j^t \leq \underline{p} \text{ for some good } j \text{ for some user } t] \leq \mathbb{P}\left[\sum_{t'=t}^{t+k-1} x_{t'j} - \sum_{t'=t}^{t+k-1} \mathbb{E}_{t'-1}[x_{t'j}] \leq -\frac{kd_j\delta}{1-\delta} \text{ for some } j, t\right],$$
$$\leq e^{\frac{-kd_j^2\delta^2}{2L^2(1-\delta)^2}},$$

for $k = o(\sqrt{n})$ which implies that the probability that the price p_j^t for some good j for some user t drops below $\underline{p}$ exponentially decays in k. Since this inequality holds for all $k = o(\sqrt{n})$, in particular, we have that the above inequality holds for $k = n^{\frac{1}{3}}$. Thus, it holds that the right hand side term goes to zero as $n \to \infty$ and so, for large n, it follows that the price of each good will always remain bounded below by $\underline{p}$ with high probability.

In particular, suppose that ϵ is the desired probability that we want to ensure that $\mathbb{P}[\sum_{t'=t}^{t+k-1} x_{t'j} \leq kd_j$ for some $j, t] \leq \epsilon$, then we require that $e^{\frac{-kd_j^2\delta^2}{2L^2(1-\delta)^2}} \leq \epsilon$, which implies that $k \geq \log(\frac{1}{\epsilon})\frac{2L^2(1-\delta)^2}{d_j^2\delta^2}$ ensures the high probability bound. In particular, this holds if $n \geq \left(\log(\frac{1}{\epsilon})\frac{2L^2(1-\delta)^2}{d_j^2\delta^2}\right)^3$.

Appendix J: Proof of Theorem 4

As with the proof of Theorem 3, the proof of Theorem 4 relies on two intermediate results. In particular, we first show in Section J.1 that if the price vector is bounded above and below by some positive constant at each iteration of Algorithm 2 with a two-stage adjustment of the step size of the price updates, then the $O(n^{2/5})$ upper bounds on both the regret and constraint violation hold. Then, in Section J.2, we modify our earlier developed potential function argument in the proof of Theorem 3 to show that if the distribution $\mathcal{D}$ satisfies Assumption 2, then the price vector $\mathbf{p}^t$ remains strictly positive and bounded throughout the operation of Algorithm 2 with a two-stage adjustment of the step size of the price updates. Note that the above two claims establish Theorem 4.

J.1. Regret and Constraint Violation Bound under Positivity and Boundedness of Prices

To prove this claim, we begin by introducing some notion. In particular, we suppose that the price vector $\mathbf{p}^t$ through the operation of Algorithm 2 with a two-stage adjustment of the price updates is such that $0 < \underline{p} \leq p_j^t \leq \bar{p}$ for all goods j and users $t \in [n]$. Furthermore, we define the $\gamma_e = O(\frac{1}{n^{2/5}})$ as the step size of the price updates in the first stage of Algorithm 2, $\gamma_p = O(\frac{1}{n^{3/5}})$ as the step size of the price updates in the second stage of Algorithm 2, and $T_e = n^{4/5}$ as the period at which the step size of Algorithm 2 changes.

Moreover, for brevity, we let $T_p = n - T_e$ be the number of periods corresponding to the second stage of Algorithm 2 for which the step size of the price updates is γ_p. Finally, we let $\mathbf{x}_1, \dots, \mathbf{x}_n$ be the allocations for the n users under the pricing policy $\boldsymbol{\pi}$ and let $\mathbf{p}^*$ be the price corresponding to the solution of the stochastic Program (37).

Then, we establish our desired regret and constraint violation bounds in three steps. First, we obtain bounds on the regret and constraint violation of Algorithm 2 in terms of the differences in the prices $\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\|_2$ and $\|\mathbf{p}^{n+1} - \mathbf{p}^*\|_2$. Then, we present upper bounds on the norm of the differences in the prices $\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\|_2$ and $\|\mathbf{p}^{n+1} - \mathbf{p}^*\|_2$. Finally, we utilize the obtained bounds on the norms of the difference in prices to obtain our desired $O(n^{2/5})$ upper bound on the regret and constraint violation.

Regret Upper Bound: We first derive an upper bound on the regret of Algorithm 2 with a two-stage adjustment of the step size of the price updates. To do so, we first recall from the generic regret bound derived in Equation (40) for any online algorithm that the expected regret:

$$\mathbb{E}[R_n(\boldsymbol{\pi})] \leq \mathbb{E}\left[\sum_{t=1}^{n}\sum_{j=1}^{m} p_j^t d_j - w_t\right] = \sum_{t=1}^{T_e} \mathbb{E}\left[\sum_{j=1}^{m} p_j^t d_j - w_t\right] + \sum_{t=T_e+1}^{n} \mathbb{E}\left[\sum_{j=1}^{m} p_j^t d_j - w_t\right] = r_e + r_p,$$

where we define $r_e = \sum_{t=1}^{T_e} \mathbb{E}\left[\sum_{j=1}^{m} p_j^t d_j - w_t\right]$ and $r_p = \sum_{t=T_e+1}^{n} \mathbb{E}\left[\sum_{j=1}^{m} p_j^t d_j - w_t\right]$.

Next, to upper bound the terms r_e and r_p, we recall that for any step size γ_t of the price updates that:

$$\left\|\mathbf{p}^{t+1}\right\|_2^2 = \left\|\mathbf{p}^t - \gamma_t (\mathbf{d} - \mathbf{x}_t)\right\|_2^2 = \left\|\mathbf{p}^t\right\|_2^2 - 2\gamma_t (\mathbf{d} - \mathbf{x}_t)^T \mathbf{p}^t + \gamma_t^2 \left\|\mathbf{d} - \mathbf{x}_t\right\|_2^2.$$

Rearranging the above equation, we obtain that:

$$(\mathbf{d} - \mathbf{x}_t)^T \mathbf{p}^t = \frac{\|\mathbf{p}^t\|_2^2 - \|\mathbf{p}^{t+1}\|_2^2}{2\gamma_t} + \frac{\gamma_t}{2}\|\mathbf{d} - \mathbf{x}_t\|_2^2.$$

Next, summing the above equation for $t \in \{1, \dots, T_e\}$, i.e., for the first-stage of Algorithm 2 for which the step size $\gamma_t = \gamma_e$ for all $t \in [T_e]$, we obtain that:

$$\begin{aligned} r_e &= \sum_{t=1}^{T_e} \mathbb{E}\left[\sum_{j=1}^{m} p_j^t d_j - w_t\right] = \sum_{t=1}^{T_e} \mathbb{E}\left[(\mathbf{d} - \mathbf{x}_t)^T \mathbf{p}^t\right] \leq \frac{\gamma_e}{2}\sum_{t=1}^{T_e} \mathbb{E}\left[\|\mathbf{d} - \mathbf{x}_t\|_2^2\right] + \mathbb{E}\left[\frac{\|\mathbf{p}^1\|_2^2 - \|\mathbf{p}^{T_e+1}\|_2^2}{2\gamma_e}\right], \\ &\overset{(a)}{\leq} \frac{\gamma_e}{2} T_e m \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}}\right)^2 + \frac{\|\mathbf{p}^1\|_2^2}{2\gamma_e}, \end{aligned} \tag{45}$$

where (a) follows by the boundedness of the consumption vector of each agent since the prices are strictly positive and bounded below by $\underline{p} > 0$.

Similarly, we can derive the following upper bound for r_p:

$$r_p = \sum_{t=T_e+1}^{n} \mathbb{E}\left[\sum_{j=1}^{m} p_j^t d_j - w_t\right] \overset{(a)}{\leq} \frac{\gamma_p}{2} T_p m \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}}\right)^2 + \frac{\mathbb{E}\left[\|\mathbf{p}^{T_e+1}\|_2^2 - \|\mathbf{p}^{n+1}\|_2^2\right]}{2\gamma_p},$$

$$
\begin{aligned}
&= \frac{\gamma_p}{2} T_p m \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}} \right)^2 + \frac{\mathbb{E}\left[(\mathbf{p}^{T_e+1} + \mathbf{p}^{T+1}) \cdot (\mathbf{p}^{T_e+1} - \mathbf{p}^{n+1}) \right]}{2\gamma_p}, \\
&\overset{(b)}{\leq} \frac{\gamma_p}{2} T_p m \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}} \right)^2 + \frac{\mathbb{E}\left[\|\mathbf{p}^{T_e+1} + \mathbf{p}^{T+1}\|_2 \, \|\mathbf{p}^{T_e+1} - \mathbf{p}^{n+1}\|_2 \right]}{2\gamma_p}, \\
&\overset{(c)}{\leq} \frac{\gamma_p}{2} T_p m \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}} \right)^2 + \frac{\bar{p}\sqrt{m}\mathbb{E}\left[\|\mathbf{p}^{T_e+1} - \mathbf{p}^{n+1}\|_2 \right]}{\gamma_p}, \\
&\overset{(d)}{\leq} \frac{\gamma_p}{2} T_p m \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}} \right)^2 + \frac{\bar{p}\sqrt{m}\mathbb{E}\left[\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\|_2 + \|\mathbf{p}^{n+1} - \mathbf{p}^*\|_2 \right]}{\gamma_p},
\end{aligned} \tag{46}
$$

where (a) follows from an analogous line of reasoning to the earlier analysis for r_e, (b) follows from the Cauchy-Schwarz inequality, (c) follows from the fact that the prices are bounded above by $\bar{p}$, and (d) follows from the triangle inequality where $\mathbf{p}^*$ is the price corresponding to the stochastic Program (37).

Summing Equations (45) and (46), we have the following upper bound on the regret:

$$
\begin{aligned}
\mathbb{E}[R_n(\boldsymbol{\pi})] &= r_e + r_p, \\
&\leq \frac{\gamma_e}{2} T_e m \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}} \right)^2 + \frac{\|\mathbf{p}^1\|_2^2}{2\gamma_e} \\
&+ \frac{\gamma_p}{2} T_p m \left(\max_{j \in [m]} d_j + \frac{\bar{w}}{\underline{p}} \right)^2 + \frac{\bar{p}\sqrt{m}\mathbb{E}\left[\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\|_2 + \|\mathbf{p}^{n+1} - \mathbf{p}^*\|_2 \right]}{\gamma_p}.
\end{aligned} \tag{47}
$$

Constraint Violation Upper Bound: Next, we provide an upper bound on the constraint violation of algorithm $\boldsymbol{\pi}$. In particular, we have that:

$$
\begin{aligned}
\mathbb{E}\left[\left\| \left(\sum_{t=1}^n \mathbf{x}_t - \mathbf{c} \right)_+ \right\|_2 \right] &\leq \mathbb{E}\left[\left\| \sum_{t=1}^n (\mathbf{x}_t - \mathbf{d}) \right\|_2 \right] = \mathbb{E}\left[\left\| \sum_{t=1}^{T_e} (\mathbf{x}_t - \mathbf{d}) + \sum_{t=T_e+1}^{T} (\mathbf{x}_t - \mathbf{d}) \right\|_2 \right], \\
&\overset{(a)}{\leq} \mathbb{E}\left[\left\| \sum_{t=1}^{T_e} \frac{\mathbf{p}^{t+1} - \mathbf{p}^t}{\gamma_e} + \sum_{t=T_e+1}^{n} \frac{\mathbf{p}^{t+1} - \mathbf{p}^t}{\gamma_p} \right\|_2 \right], \\
&= \mathbb{E}\left[\left\| \frac{\mathbf{p}^{T_e+1} - \mathbf{p}^1}{\gamma_e} + \frac{\mathbf{p}^{n+1} - \mathbf{p}^{T_e+1}}{\gamma_p} \right\|_2 \right], \\
&\overset{(b)}{\leq} \mathbb{E}\left[\left\| \frac{\mathbf{p}^{T_e+1} - \mathbf{p}^1}{\gamma_e} \right\|_2 + \left\| \frac{\mathbf{p}^{n+1} - \mathbf{p}^{T_e+1}}{\gamma_p} \right\|_2 \right], \\
&\leq \frac{1}{\gamma_e} \mathbb{E}\left[\|\mathbf{p}^{T_e+1} + \mathbf{p}^1\|_2 \right] + \frac{1}{\gamma_p} \mathbb{E}\left[\|\mathbf{p}^{n+1} - \mathbf{p}^{T_e+1}\|_2 \right], \\
&\overset{(c)}{\leq} \frac{2\bar{p}\sqrt{m}}{\gamma_e} + \frac{1}{\gamma_p} \mathbb{E}\left[\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\|_2 + \|\mathbf{p}^{n+1} - \mathbf{p}^*\|_2 \right],
\end{aligned} \tag{48}
$$

where (a) follows from the price update rule of Algorithm 2 and (b) and (c) follow by the triangle inequality.

Bound on Regret and Constraint Violation: Combining the above obtained relations in Equations (47) and (48), we obtain the following bound on the sum of the regret and constraint violation of Algorithm 2 with a two-stage adjustment in the step size:

$$
\mathbb{E}[R_n(\boldsymbol{\pi}) + V_n(\boldsymbol{\pi})] \leq \frac{2\bar{p}\sqrt{m}}{\gamma_e} + \frac{1}{\gamma_p} \mathbb{E}\left[\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\|_2 + \|\mathbf{p}^{n+1} - \mathbf{p}^*\|_2 \right] \tag{49}
$$

$$+\frac{\gamma_e}{2}T_e m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2+\frac{\|\mathbf{p}^1\|_2^2}{2\gamma_e}$$
$$+\frac{\gamma_p}{2}T_p m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2+\frac{\bar{p}\sqrt{m}\mathbb{E}\left[\|\mathbf{p}^{T_e+1}-\mathbf{p}^*\|_2+\|\mathbf{p}^{n+1}-\mathbf{p}^*\|_2\right]}{\gamma_p}. \tag{50}$$

In the above upper bound, we combine the terms containing γ_e and those containing γ_p to obtain two relations:

$$W_e=\frac{2\bar{p}\sqrt{m}}{\gamma_e}+\frac{\gamma_e}{2}T_e m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2+\frac{\|\mathbf{p}^1\|_2^2}{2\gamma_e}=\frac{\gamma_e}{2}T_e m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2+\frac{4\bar{p}\sqrt{m}+\|\mathbf{p}^1\|_2^2}{2\gamma_e}.$$

$$W_p=\frac{\gamma_p}{2}T_p m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2+\frac{\bar{p}\sqrt{m}+1}{\gamma_p}\mathbb{E}\left[\left\|\mathbf{p}^{T_e+1}-\mathbf{p}^*\right\|_2+\left\|\mathbf{p}^{n+1}-\mathbf{p}^*\right\|_2\right].$$

From the above relation for W_e, note that setting $\gamma_e=\frac{1}{\sqrt{T_e}}$, we obtain that $W_e=O(\sqrt{T_e})=O(n^{2/5})$.

Bound on Difference in Prices: To upper bound W_p, we now present bounds on the terms $\mathbb{E}\left[\|\mathbf{p}^{T_e+1}-\mathbf{p}^*\|_2\right]$ and $\mathbb{E}\left[\|\mathbf{p}^{n+1}-\mathbf{p}^*\|_2\right]$. To this end, consider the following upper bound for the square of the first term:

$$\begin{aligned}
\mathbb{E}\left[\left\|\mathbf{p}^{T_e+1}-\mathbf{p}^*\right\|_2^2\right]&=\mathbb{E}\left[\left\|\mathbf{p}^{T_e}-\gamma_e(\mathbf{d}-\mathbf{x}_{T_e})-\mathbf{p}^*\right\|_2^2\right],\\
&=\mathbb{E}\left[\left\|\mathbf{p}^{T_e}-\mathbf{p}^*\right\|_2^2+\gamma_e^2\left\|(\mathbf{d}-\mathbf{x}_{T_e})\right\|_2^2-2\gamma_e(\mathbf{d}-\mathbf{x}_{T_e})^T(\mathbf{p}^{T_e}-\mathbf{p}^*)\right],\\
&\overset{(a)}{\leq}\mathbb{E}\left[\left\|\mathbf{p}^{T_e}-\mathbf{p}^*\right\|_2^2\right]+\gamma_e^2 m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2-2\gamma_e\mathbb{E}\left[(\mathbf{d}-\mathbf{x}_{T_e})^T(\mathbf{p}^{T_e}-\mathbf{p}^*)\right],\\
&\overset{(b)}{\leq}\mathbb{E}\left[\left\|\mathbf{p}^{T_e}-\mathbf{p}^*\right\|_2^2\right]+\gamma_e^2 m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2-2\gamma_e\mathbb{E}\left[(D(\mathbf{p}^{T_e})-D(\mathbf{p}^*))\right],
\end{aligned}$$

where (a) follows by the boundedness of the consumption vector as the prices are strictly positive and lower bounded by $\underline{p}>0$ through the course of the Algorithm 2 and (b) follows by conditioning on the history and taking the expectation and by the convexity of the dual function $D(\mathbf{p})$ (see Equation (26) and recall that $\mathbb{E}[\mathbf{d}-\mathbf{x}_{T_e}]$ is a sub-gradient of the dual function given that the price $\mathbf{p}^{T_e}$ for user T_e as noted in Section 5.2).

Next, to upper bound the right hand side of the above term, as in the proof of Theorem 2, we reparametrize the dual objective by introducing a variable $\alpha_k=\min_j\frac{p_j}{\tilde{u}_{kj}}$ for all utility vectors $\tilde{\mathbf{u}}_k$ drawn from the discrete distribution $\mathcal{D}$ with finite support. Then, from our earlier analysis in Equations (27) and (29), we have that:

$$\begin{aligned}
\mathbb{E}\left[\left\|\mathbf{p}^{T_e+1}-\mathbf{p}^*\right\|_2^2\right]&\overset{(c)}{\leq}\mathbb{E}\left[\left\|\mathbf{p}^{T_e}-\mathbf{p}^*\right\|_2^2\right]+\gamma_e^2 m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2-2\gamma_e\mathbb{E}\left[(D(\boldsymbol{\alpha}^{T_e})-D(\boldsymbol{\alpha}^*))\right],\\
&\overset{(d)}{\leq}\mathbb{E}\left[\left\|\mathbf{p}^{T_e}-\mathbf{p}^*\right\|_2^2\right]+\gamma_e^2 m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2-2\gamma_e\eta\mathbb{E}\left[\left\|\boldsymbol{\alpha}^{T_e}-\boldsymbol{\alpha}^*\right\|_2^2\right],\\
&\overset{(e)}{\leq}\mathbb{E}\left[\left\|\mathbf{p}^{T_e}-\mathbf{p}^*\right\|_2^2\right]+\gamma_e^2 m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2-\frac{2\gamma_e\eta}{\sqrt{m}\bar{u}}\mathbb{E}\left[\left\|\mathbf{p}^{T_e}-\mathbf{p}^*\right\|_2^2\right],\\
&=\left(1-\frac{2\gamma_e\eta}{\sqrt{m}\bar{u}}\right)\mathbb{E}\left[\left\|\mathbf{p}^{T_e}-\mathbf{p}^*\right\|_2^2\right]+\gamma_e^2 m\left(\max_{j\in[m]} d_j+\frac{\bar{w}}{\underline{p}}\right)^2,
\end{aligned}$$

$$\leq \left(1 - \frac{2\gamma_e \eta}{\sqrt{m}\bar{u}}\right)^{T_e} \left\|\mathbf{p}^1 - \mathbf{p}^*\right\|_2^2 + \sum_{l=0}^{T_e-1} \gamma_e^2 m \left(\max_{j\in[m]} d_j + \frac{\bar{w}}{\underline{p}}\right)^2 \left(1 - \frac{2\gamma_e \eta}{\sqrt{m}\bar{u}}\right)^l,$$

where (c) follows from the variable transformation $\alpha_k = \min_j \frac{p_j}{\bar{u}_{kj}}$ as in the proof of Lemma 3, (d) follows from Equation (27) and (e) follows from Equation (29), derived in our analysis in the proof of Lemma 3 for discrete distributions with finite support.

Next, defining $\zeta_1 = \frac{2\eta}{\sqrt{m}\bar{u}}$ and $\zeta_2 = m\left(\max_{j\in[m]} d_j + \frac{\bar{w}}{\underline{p}}\right)^2$, we have from the above relations that:

$$\begin{aligned}
\mathbb{E}\left[\left\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\right\|_2^2\right] &\leq (1-\gamma_e\zeta_1)^{T_e} \left\|\mathbf{p}^1 - \mathbf{p}^*\right\|_2^2 + \gamma_e^2 \zeta_2 \sum_{l=0}^{T_e-1} (1-\gamma_e\zeta_1)^l, \\
&\overset{(a)}{=} (1-\gamma_e\zeta_1)^{T_e} \left\|\mathbf{p}^1 - \mathbf{p}^*\right\|_2^2 + \gamma_e^2\zeta_2 \frac{1-(1-\gamma_e\zeta_1)^{T_e}}{\gamma_e\zeta_1}, \\
&\overset{(b)}{\leq} (1-\gamma_e\zeta_1)^{T_e} \left\|\mathbf{p}^1 - \mathbf{p}^*\right\|_2^2 + \frac{\gamma_e\zeta_2}{\zeta_1}, \\
&\overset{(c)}{\leq} \frac{\left\|\mathbf{p}^1 - \mathbf{p}^*\right\|_2^2}{1+\gamma_e\zeta_1 T_e} + \frac{\gamma_e\zeta_2}{\zeta_1}, \\
&\leq \frac{\left\|\mathbf{p}^1 - \mathbf{p}^*\right\|_2^2}{\gamma_e\zeta_1 T_e} + \frac{\gamma_e\zeta_2}{\zeta_1},
\end{aligned} \tag{51}$$

where (a) follows from the formula of the sum of a geometric series, (b) follows as $1-\gamma_e\zeta_1 \geq 0$, as we can select $\gamma_e = \frac{D_e}{n^{2/5}}$ such that the constant $D_e \leq \frac{1}{\zeta_1}$, and (c) follows as $(1-x)^r \leq \frac{1}{1+rx}$ for $x \in (0,1)$.

Following the above analysis, we can analogously derive that:

$$\mathbb{E}\left[\left\|\mathbf{p}^{n+1} - \mathbf{p}^*\right\|_2^2\right] \leq \frac{\mathbb{E}[\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\|_2^2]}{\gamma_p\zeta_1 T_p} + \frac{\gamma_p\zeta_2}{\zeta_1}. \tag{52}$$

Bound on W_p: We now use Equations (51) and (52) to compute the bound for V_p. In particular, we get that:

$$\begin{aligned}
W_p &= \frac{\gamma_p}{2} T_p \zeta_2 + \frac{\bar{p}\sqrt{m}+1}{\gamma_p} \mathbb{E}\left[\left\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\right\|_2 + \left\|\mathbf{p}^{n+1} - \mathbf{p}^*\right\|_2\right], \\
&\overset{(a)}{\leq} \frac{\gamma_p}{2} T_p \zeta_2 + \frac{\bar{p}\sqrt{m}+1}{\gamma_p} \left[\sqrt{\mathbb{E}\left[\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\|_2^2\right]} + \sqrt{\mathbb{E}\left[\|\mathbf{p}^{n+1} - \mathbf{p}^*\|_2^2\right]}\right] \\
&\overset{(b)}{\leq} \frac{\gamma_p}{2} T_p \zeta_2 + \frac{\bar{p}\sqrt{m}+1}{\gamma_p} \left(\sqrt{\frac{\|\mathbf{p}^1 - \mathbf{p}^*\|_2^2}{\gamma_e\zeta_1 T_e} + \frac{\gamma_e\zeta_2}{\zeta_1}} + \sqrt{\frac{\mathbb{E}[\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\|_2^2]}{\gamma_p\zeta_1 T_p} + \frac{\gamma_p\zeta_2}{\zeta_1}}\right), \\
&\overset{(c)}{\leq} \frac{\gamma_p}{2} T_p \zeta_2 + \frac{\bar{p}\sqrt{m}+1}{\gamma_p} \left(\sqrt{\frac{\|\mathbf{p}^1 - \mathbf{p}^*\|_2^2}{\gamma_e\zeta_1 T_e}} + \sqrt{\frac{\gamma_e\zeta_2}{\zeta_1}} + \sqrt{\frac{\mathbb{E}[\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\|_2^2]}{\gamma_p\zeta_1 T_p}} + \sqrt{\frac{\gamma_p\zeta_2}{\zeta_1}}\right), \\
&\overset{(d)}{=} O(\gamma_p T_p + \gamma_p^{-1}(\gamma_e^{-1/2} T_e^{-1/2} + \gamma_p^{-1/2} T_p^{-1/2} + \gamma_e^{1/2} + \gamma_p^{1/2}), \\
&= O(\gamma_p T_p + \gamma_p^{-1}\gamma_e^{-1/2} T_e^{-1/2} + \gamma_p^{-3/2} T_p^{-1/2} + \gamma_p^{-1}\gamma_e^{1/2} + \gamma_p^{-1/2}), \\
&\overset{(e)}{=} O(n^{2/5}),
\end{aligned}$$

where (a) follows by Jensen's inequality, (b) follows by our derived relations in Equations (51) and (52), (c) follows as $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$, (d) follows by dropping all the constants, and (e) follows by plugging in the expressions for γ_p, γ_e, T_p, and T_e and simplifying.

$O(n^{2/5})$ *Bound:* Thus, we have shown in our above analysis that both W_e and W_p are upper bounded by $O(n^{2/5})$. Consequently, from Equation (49) we obtain the following bound on the regret and constraint violation of Algorithm 2 with a two-stage adjustment in the step size of the price updates:

$$\mathbb{E}[R_n(\boldsymbol{\pi}) + V_n(\boldsymbol{\pi})] \leq W_e + W_p \leq O(n^{2/5}),$$

which implies that both the regret and constraint violation of Algorithm 2 with a two-stage adjustment in the step size of the price updates are bounded by $O(n^{2/5})$, establishing our desired result.

J.2. Positivity and Boundedness of Prices

We note that the proof of the positivity and boundedness of prices throughout the operation of Algorithm 2 with a two-stage adjustment in the step size of the price updates follows almost entirely analogously to the proof of Lemma 7 for the case of Algorithm 2 with a fixed step size. To see this, consider the two stages corresponding to the Algorithm 2 when the step size is fixed to γ_e and γ_p, respectively. In this case, note that the step size is fixed until the arrival of the first T_e users. Consequently, following an entirely analogous line of reasoning to the proof of Lemma 7, the price vectors $\mathbf{p}^1, \mathbf{p}^2, \ldots, \mathbf{p}^{T_e}, \mathbf{p}^{T_e+1}$ remain strictly positive and bounded as the step size is fixed to γ_e across the first T_e periods. Next, taking the price vector $\mathbf{p}^{T_e+1}$ as the initial price vector for the second stage, which is strictly positive and bounded as noted above, it also follows from an entirely analogous line of reasoning to the proof of Lemma 7 that the price vectors $\mathbf{p}^{T_e+2}, \mathbf{p}^{T_e+3}, \ldots, \mathbf{p}^n$, i.e., the prices in the second stage of Algorithm 2, remain strictly positive and bounded as the step size is fixed to γ_p for the remaining periods. Hence, from an almost entirely analogous line of reasoning to the proof of Lemma 7, it follows that the prices are strictly positive and bounded throughout the operation of Algorithm 2 with a two-stage adjustment of the step size of the price updates.

Appendix K: Proofs of Results in Section 6

K.1. Proof of Theorem 5

To prove this result, we begin by defining $U_{\tau^{\boldsymbol{\pi}}+1:n}$ as the budget weighted logarithmic utility objective ranging from the periods $\tau^{\boldsymbol{\pi}} + 1$ to n. Then, the expected regret for algorithm $\boldsymbol{\pi}^f$ can be expressed as follows:

$$\begin{aligned}
R_n(\boldsymbol{\pi}^f) &= \mathbb{E}[U_n^* - U_n(\boldsymbol{\pi}^f)], \\
&= \mathbb{E}[U_{\tau^{\boldsymbol{\pi}}}^* - U_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}^f)] + \mathbb{E}[U_{\tau^{\boldsymbol{\pi}}+1:n}^* - U_{\tau^{\boldsymbol{\pi}}+1:n}(\boldsymbol{\pi}^f)], \\
&= R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) + \mathbb{E}[U_{\tau^{\boldsymbol{\pi}}+1:n}^* - U_{\tau^{\boldsymbol{\pi}}+1:n}(\boldsymbol{\pi}^f)],
\end{aligned} \tag{53}$$

where the final equality follows as the algorithms $\boldsymbol{\pi}$ and $\boldsymbol{\pi}^f$ are identical up until period $\tau^{\boldsymbol{\pi}}$. Then, to establish the desired upper bound on the regret, we now bound second term on the right hand side of the above equation.

Regret from period $\tau^{\boldsymbol{\pi}}+1$ to n: Recall that users from periods $\tau^{\boldsymbol{\pi}}+1$ through n are provided $\frac{\epsilon}{n}$ units of each resource. Thus, it holds that

$$\begin{aligned}
\mathbb{E}[U^*_{\tau^{\boldsymbol{\pi}}+1:n} - U_{\tau^{\boldsymbol{\pi}}+1:n}(\boldsymbol{\pi}^f)] &= \mathbb{E}\left[\sum_{t=\tau^{\boldsymbol{\pi}}+1}^{n} w_t \log(\sum_{j=1}^{m} u_{tj} x^*_{tj})\right] - \mathbb{E}\left[\sum_{t=\tau^{\boldsymbol{\pi}}+1}^{n} w_t \log(\sum_{j=1}^{m} u_{tj} x_{tj})\right], \\
&\overset{(a)}{\leq} \mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]\bar{w} \log\left(\sum_{j=1}^{m} \bar{u}\frac{\bar{w}}{\underline{p}}\right) - \mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]\left[\underline{w}\log\left(\underline{u}\frac{\epsilon}{n}\right)\right], \\
&= \mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]\left[\bar{w}\log\left(m\bar{u}\frac{\bar{w}}{\underline{p}}\right) - \underline{w}\log(\underline{u}\epsilon) + \underline{w}\log(n)\right], \\
&= O(\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]\log(n)), \qquad (54)
\end{aligned}$$

where (a) follows as $\bar{w}$, $\underline{w}$, $\bar{u}$, $\underline{u}$ are strictly positive upper and lower bound on the budgets and utilities, respectively, and the fact that all users have at least one resource for which $u_{tj} \geq \underline{u} > 0$. Furthermore, $\underline{p} > 0$ corresponds to a lower bound on the optimal price at any period, which holds under Assumption 1.

Final Regret Bound for Algorithms 1 and 2: Combining the above obtained relationship for the second component of the regret on the right hand side of Equation (53), we obtain from Equation (54) the following upper bound on the regret:

$$\begin{aligned}
R_n(\boldsymbol{\pi}^f) &= R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) + \mathbb{E}[U^*_{\tau^{\boldsymbol{\pi}}+1:n} - U_{\tau^{\boldsymbol{\pi}}+1:n}(\boldsymbol{\pi}^f)], \\
&\leq R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) + O(\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]\log(n)),
\end{aligned}$$

which establishes our result.

K.2. Proof of Corollary 2

In this proof, for brevity, we use Algorithm 2 to refer to the revealed preference algorithm with a fixed step size of $O(\frac{1}{\sqrt{n}})$.

To prove our desired regret bound, we first note from Theorem 5 that we have the following upper bound on the regret of the feasible variant $\boldsymbol{\pi}^f$ of Algorithm 2:

$$R_n(\boldsymbol{\pi}^f) \leq R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) + O(\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]\log(n)), \qquad (55)$$

where $\boldsymbol{\pi}$ corresponds to Algorithm 2. To bound the terms on the right hand side of the above equation, first note by the analysis in the proof of Theorem 3 that $R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) \leq O(\sqrt{\tau^{\boldsymbol{\pi}}})$. Next, we obtain an upper bound on the second term on the right hand side of the above equation by obtaining a bound on $\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]$.

Upper Bound on $\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]$: To bound $\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]$, first note by the definition of $\tau^{\boldsymbol{\pi}}$ that there is some good $j \in [m]$ for which the following inequality holds:

$$\sum_{t=1}^{\tau^{\boldsymbol{\pi}}} x_{tj} + \frac{\bar{w}}{\underline{p}} \geq c_j = nd_j,$$

where $\underline{p} > 0$ corresponds to a lower bound on the price in the update step of Algorithm 2 (see Lemma 7). Subtracting $\tau^{\boldsymbol{\pi}} d_j$ to both sides of the above inequality and rearranging, we obtain the following upper bound on $n - \tau^{\boldsymbol{\pi}}$:

$$\tau^{\boldsymbol{\pi}} d_j - \sum_{t=1}^{\tau^{\boldsymbol{\pi}}} x_{tj} \leq \tau^{\boldsymbol{\pi}} d_j - d_j n + \frac{\bar{w}}{\underline{p}},$$

$$\implies (n - \tau^{\boldsymbol{\pi}}) \leq \frac{\frac{\bar{w}}{\underline{p}} - (\tau^{\boldsymbol{\pi}} d_j - \sum_{t=1}^{\tau^{\boldsymbol{\pi}}} x_{tj})}{d_j} \tag{56}$$

Next, to bound the right hand side of the above equation, we note by the price update rule in Algorithm 2 that:

$$\sum_{t=1}^{\tau^{\boldsymbol{\pi}}} -(d_j - x_{tj}) = \frac{1}{\gamma} \sum_{t=1}^{\tau^{\boldsymbol{\pi}}} (p_j^{t+1} - p_j^t) = \frac{1}{\gamma}(p_j^{\tau^{\boldsymbol{\pi}}+1} - p_j^1) \leq \frac{1}{\gamma}(\bar{p} - p_j^1),$$

where recall that $\bar{p}$ represents an upper bound on the price at any stage of Algorithm 2 as established in the proof of Theorem 3. Then, combining the above inequality with Equation (56), we obtain that:

$$n - \tau^{\boldsymbol{\pi}} \leq \frac{\frac{\bar{w}}{\underline{p}} + \frac{1}{\gamma}(\bar{p} - p_j^1)}{d_j} = O(\sqrt{n}). \tag{57}$$

Final Regret Bound: Combining the above obtained relationship for $\mathbb{E}[(n - \tau^{\boldsymbol{\pi}})]$, we obtain the following upper bound on the regret of $\boldsymbol{\pi}^f$:

$$R_n(\boldsymbol{\pi}^f) \leq R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) + O(\mathbb{E}[(n - \tau^{\boldsymbol{\pi}})] \log(n)) \overset{(a)}{\leq} O(\sqrt{\tau^{\boldsymbol{\pi}}}) + O(\sqrt{n} \log(n)),$$
$$\overset{(b)}{\leq} O(\sqrt{n}) + O(\sqrt{n} \log(n)) = O(\sqrt{n} \log(n)),$$

where (a) follows from Equation (57) and the fact that $R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) \leq O(\sqrt{\tau^{\boldsymbol{\pi}}})$ from the analysis in Theorem 3. Furthermore, (b) follows from the fact that $\tau^{\boldsymbol{\pi}} \leq n$. The above inequality establishes our desired result that $R_n(\boldsymbol{\pi}^f) \leq O(\sqrt{n} \log(n))$.

K.3. Proof of Corollary 3

In this proof, for brevity, we use Algorithm 2 to refer to the revealed preference algorithm with a two-stage adjustment of the step size.

We first note From Theorem 5 that we have the following upper bound on the regret of the feasible variant $\boldsymbol{\pi}^f$ of Algorithm 2:

$$R_n(\boldsymbol{\pi}^f) \leq R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) + O(\mathbb{E}[(n - \tau^{\boldsymbol{\pi}})] \log(n)), \tag{58}$$

where $\boldsymbol{\pi}$ corresponds to Algorithm 2. To bound the terms on the right hand side of the above equation, first note by the analysis in the proof of Theorem 4 that $R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) \leq O((\tau^{\boldsymbol{\pi}})^{2/5})$. Next, we obtain an upper bound on the second term on the right hand side of the above equation by obtaining a bound on $\mathbb{E}[(n - \tau^{\boldsymbol{\pi}})]$.

Upper Bound on $\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]$: To bound $\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]$, first note by the definition of $\tau^{\boldsymbol{\pi}}$ that there is some good $j \in [m]$ for which the following inequality holds:

$$\sum_{t=1}^{\tau^{\boldsymbol{\pi}}} x_{tj} + \frac{\bar{w}}{\underline{p}} \geq c_j = nd_j, \tag{59}$$

where $\underline{p} > 0$ corresponds to a lower bound on the price in the update step of Algorithm 2 (see Lemma 7). The above inequality implies that:

$$\frac{\tau^{\boldsymbol{\pi}}\bar{w}}{\underline{p}} + \frac{\bar{w}}{\underline{p}} \geq nd_j \implies \tau^{\boldsymbol{\pi}} \geq n\frac{d_j\underline{p}}{\bar{w}} - 1.$$

Then, under the condition that $n \geq \max\{1, \left(\frac{2\bar{w}}{\underline{p}\min_j d_j}\right)^5\}$, using the above inequality, it is straightforward to check that $\tau^{\boldsymbol{\pi}} - T_e = \tau^{\boldsymbol{\pi}} - n^{4/5} \geq n\frac{d_j\underline{p}}{\bar{w}} - 1 - n^{4/5} \geq 0$, i.e., $\tau^{\boldsymbol{\pi}} - T_e \geq \Omega(n)$ for all $n \geq \max\{1, \left(\frac{2\bar{w}}{\underline{p}\min_j d_j}\right)^5\}$, where recall from the proof of Theorem 4 that $T_e = n^{4/5}$ represents the period at which the step size of Algorithm 2 is adjusted.

Next, subtracting $\tau^{\boldsymbol{\pi}} d_j$ on both sides of the inequality in Equation (59) and rearranging, we obtain the following upper bound on $n - \tau^{\boldsymbol{\pi}}$:

$$\begin{aligned}
\tau^{\boldsymbol{\pi}} d_j - \sum_{t=1}^{\tau^{\boldsymbol{\pi}}} x_{tj} &\leq \tau^{\boldsymbol{\pi}} d_j - d_j T + \frac{\bar{w}}{\underline{p}}, \\
\implies (n - \tau^{\boldsymbol{\pi}}) &\leq \frac{\frac{\bar{w}}{\underline{p}} - (\tau^{\boldsymbol{\pi}} d_j - \sum_{t=1}^{\tau^{\boldsymbol{\pi}}} x_{tj})}{d_j}
\end{aligned} \tag{60}$$

Next, to bound the right hand side of the above equation, we note by the price update rule in Algorithm 2 that:

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{\tau^{\boldsymbol{\pi}}} -(d_j - x_{tj})\right] &= \mathbb{E}\left[\frac{1}{\gamma_e}\sum_{t=1}^{T_e}(p_j^{t+1} - p_j^t) + \frac{1}{\gamma_p}\sum_{t=T_e+1}^{\tau^{\boldsymbol{\pi}}}(p_j^{t+1} - p_j^t)\right], \\
&= \frac{1}{\gamma_e}\mathbb{E}\left[(p_j^{T_e+1} - p_j^1)\right] + \frac{1}{\gamma_p}\mathbb{E}\left[(p_j^{\tau^{\boldsymbol{\pi}}+1} - p_j^{T_e+1})\right], \\
&\overset{(a)}{\leq} \frac{\bar{p}}{\gamma_e} + \frac{1}{\gamma_p}\mathbb{E}\left[|p_j^{\tau^{\boldsymbol{\pi}}+1} - p_j^{T_e+1}|\right], \\
&\leq \frac{\bar{p}}{\gamma_e} + \frac{1}{\gamma_p}\mathbb{E}\left[\left\|\mathbf{p}^{\tau^{\boldsymbol{\pi}}+1} - \mathbf{p}^{T_e+1}\right\|_2\right], \\
&\overset{(b)}{\leq} \frac{\bar{p}}{\gamma_e} + \frac{1}{\gamma_p}\left[\mathbb{E}\left[\left\|\mathbf{p}^{\tau^{\boldsymbol{\pi}}+1} - \mathbf{p}^*\right\|_2\right] + \mathbb{E}\left[\left\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\right\|_2\right]\right],
\end{aligned} \tag{61}$$

where (a) follows as $\bar{p}$ represents an upper bound on the price at any stage of Algorithm 2 and (b) follows by the triangle inequality.

From our earlier analysis (see Equations (51) and (52)) in the proof of Theorem 4, we have that the terms $\mathbb{E}\left[\left\|\mathbf{p}^{\tau^{\boldsymbol{\pi}}+1} - \mathbf{p}^*\right\|_2\right]$ and $\mathbb{E}\left[\left\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\right\|_2\right]$ on the right hand side of the above inequality can be bounded as follows:

$$\mathbb{E}\left[\left\|\mathbf{p}^{T_e+1} - \mathbf{p}^*\right\|_2^2\right] \leq \frac{\left\|\mathbf{p}^1 - \mathbf{p}^*\right\|_2^2}{\gamma_e\zeta_1 T_e} + \frac{\gamma_e\zeta_2}{\zeta_1} \tag{62}$$

$$\mathbb{E}\left[\left\|\mathbf{p}^{\tau^{\boldsymbol{\pi}}}-\mathbf{p}^*\right\|_2^2\right] \leq \frac{\mathbb{E}\left[\|\mathbf{p}^{T_e+1}-\mathbf{p}^*\|_2^2\right]}{\gamma_p\zeta_1(\tau^{\boldsymbol{\pi}}-T_e)}+\frac{\gamma_p\zeta_2}{\zeta_1}, \tag{63}$$

where recall that ζ_1, ζ_2 are constants, $\gamma_e = O(\frac{1}{n^{2/5}})$ is the step size of the price updates in the first stage of Algorithm 2, $\gamma_p = O(\frac{1}{n^{3/5}})$ is the step size of the price updates in the second stage of Algorithm 2, and T_e is the period at which the step size of Algorithm 2 changes.

Finally, plugging in Equations (62) and (63) into the right hand side of Equation (61), we get:

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{\tau^{\boldsymbol{\pi}}} -(d_j - x_{tj})\right] &\leq \frac{\bar{p}}{\gamma_e} + \frac{1}{\gamma_p}\left[\mathbb{E}\left[\left\|\mathbf{p}^{\tau^{\boldsymbol{\pi}}+1}-\mathbf{p}^*\right\|_2\right] + \mathbb{E}\left[\|\mathbf{p}^{T_e+1}-\mathbf{p}^*\|_2\right]\right], \\
&\overset{(a)}{\leq} \frac{\bar{p}}{\gamma_e} + \frac{1}{\gamma_p}\left[\sqrt{\frac{\|\mathbf{p}^1-\mathbf{p}^*\|_2^2}{\gamma_e\zeta_1 T_e}+\frac{\gamma_e\zeta_2}{\zeta_1}} + \sqrt{\frac{\mathbb{E}\left[\|\mathbf{p}^{T_e+1}-\mathbf{p}^*\|_2^2\right]}{\gamma_p\zeta_1(\tau^{\boldsymbol{\pi}}-T_e)}+\frac{\gamma_p\zeta_2}{\zeta_1}}\right], \\
&\overset{(b)}{\leq} \frac{\bar{p}}{\gamma_e} + \frac{1}{\gamma_p}\left[\sqrt{\frac{\|\mathbf{p}^1-\mathbf{p}^*\|_2^2}{\gamma_e\zeta_1 T_e}}+\sqrt{\frac{\gamma_e\zeta_2}{\zeta_1}} + \sqrt{\frac{\mathbb{E}\left[\|\mathbf{p}^{T_e+1}-\mathbf{p}^*\|_2^2\right]}{\gamma_p\zeta_1(\tau^{\boldsymbol{\pi}}-T_e)}}+\sqrt{\frac{\gamma_p\zeta_2}{\zeta_1}}\right], \\
&\overset{(c)}{=} O(\gamma_e^{-1}+\gamma_p^{-1}(\gamma_e^{-1/2}T_e^{-1/2}+\gamma_p^{-1/2}(\tau^{\boldsymbol{\pi}}-T_e)^{-1/2}+\gamma_e^{1/2}+\gamma_p^{1/2}), \\
&\overset{(d)}{=} O(\gamma_e^{-1}+\gamma_p^{-1}\gamma_e^{-1/2}T_e^{-1/2}+\gamma_p^{-3/2}n^{-1/2}+\gamma_p^{-1}\gamma_e^{1/2}+\gamma_p^{-1/2}), \\
&\overset{(e)}{=} O(n^{2/5}),
\end{aligned}$$

where (a) follows by Cauchy-Schwarz inequality and plugging in Equations (62) and (63) into the right hand side of Equation (61), (b) follows as $\sqrt{a+b} \leq \sqrt{a}+\sqrt{b}$, and (c) follows by dropping the constants and only retaining the terms that depend on the number of users n. Moreover, (d) follows as $\tau^{\boldsymbol{\pi}} - T_e = O(n)$, as $\tau^{\boldsymbol{\pi}} - T_e \geq \Omega(n)$ as noted earlier when $n \geq \max\{1, \left(\frac{2\bar{w}}{\underline{p}\min_j d_j}\right)^5\}$ and $T_e = O(n^{4/5})$. Finally, (e) follows by plugging in the expressions for γ_p, γ_e, and T_e (see proof of Theorem 4) and simplifying.

Then, combining the above inequality with Equation (60), we obtain that:

$$\mathbb{E}\left[n-\tau^{\boldsymbol{\pi}}\right] \leq \frac{\frac{\bar{w}}{\underline{p}} - \left(\mathbb{E}\left[\tau^{\boldsymbol{\pi}} d_j - \sum_{t=1}^{\tau^{\boldsymbol{\pi}}} x_{tj}\right]\right)}{d_j} = O(n^{2/5}). \tag{64}$$

Final Regret Bound: Combining the above obtained relationship for $\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]$, we obtain the following upper bound on the regret of $\boldsymbol{\pi}^f$:

$$\begin{aligned}
R_n(\boldsymbol{\pi}^f) \leq R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) + O(\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]\log(n)) &\overset{(a)}{\leq} O((\tau^{\boldsymbol{\pi}})^{2/5}) + O(n^{2/5}\log(n)), \\
&\overset{(b)}{\leq} O(n^{2/5}) + O(n^{2/5}\log(n)) = O(n^{2/5}\log(n)),
\end{aligned}$$

where (a) follows from Equation (64) and the fact that $R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) \leq O((\tau^{\boldsymbol{\pi}})^{2/5})$ from the analysis in Theorem 4. Furthermore, (b) follows from the fact that $\tau^{\boldsymbol{\pi}} \leq n$. The above inequality establishes our desired result that $R_n(\boldsymbol{\pi}^f) \leq O(n^{2/5}\log(n))$.

K.4. Proof of Corollary 4

From Theorem 5, we have the following upper bound on the regret of the feasible variant of Algorithm 1:

$$R_n(\boldsymbol{\pi}^f) \leq R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) + O(\mathbb{E}[(n - \tau^{\boldsymbol{\pi}})] \log(n)), \tag{65}$$

where $\boldsymbol{\pi}$ corresponds to Algorithm 1. To bound the terms on the right hand side of the equation, we first recall that $O(\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]$ is constant for Algorithm 1 (see Equation (25) in the proof of Lemma 2). Next, following our analysis in the proof of Theorem 2, note that $R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) \leq O(\log(\tau^{\boldsymbol{\pi}}))$. Thus, we get the following upper bound on the regret of $\boldsymbol{\pi}^f$:

$$R_n(\boldsymbol{\pi}^f) \leq R_{\tau^{\boldsymbol{\pi}}}(\boldsymbol{\pi}) + O(\mathbb{E}[(n - \tau^{\boldsymbol{\pi}})] \log(n)) \overset{(a)}{\leq} O(\log(\tau^{\boldsymbol{\pi}})) + O(\log(n)) \overset{(b)}{\leq} O(\log(n)),$$

where (a) follows as $\mathbb{E}[(n-\tau^{\boldsymbol{\pi}})]$ is a constant for Algorithm 1 as noted above, and (b) follows as $\tau^{\boldsymbol{\pi}} \leq n$. The above inequality establishes our desired result that $R_n(\boldsymbol{\pi}^f) \leq O(\log(n))$.

Appendix L: Additional Details on Benchmarks

In this section, we provide more details on two of the benchmarks (i.e., the *Stochastic Program* and *Dynamic Learning SAA* benchmarks) to which we compare our revealed preference algorithms in our experiments. Both these benchmarks are akin to several classical algorithms developed in the online resource allocation literature (Li and Ye 2022, Agrawal et al. 2014) and assume access to additional information on users' utility and budget parameters. In particular, the first benchmark assumes knowledge of the distribution $\mathcal{D}$ from which the budget and utility parameters are drawn, as is the case for an algorithm that sets expected equilibrium prices. The second benchmark assumes that users' utility and budget parameters are revealed to the central planner when they enter the market and can be used to set prices for subsequent users. We mention that these algorithms are solely for benchmark purposes, and thus we do not discuss the practicality of the corresponding informational assumptions of these benchmarks. We also reiterate that, as opposed to these benchmarks, the price updates in Algorithm 2 only rely on users' revealed preferences rather than relying on additional information on their budget and utility parameters.

Stochastic Program: We begin with the benchmark wherein the distribution $\mathcal{D}$ from which the budget and utility parameters are generated i.i.d. is known. In this case, the SAA Problem (5) is related to the following stochastic program

$$\min_{\mathbf{p}} \quad D(\mathbf{p}) = \sum_{j=1}^{m} p_j d_j + \mathbb{E}_{(w,\mathbf{u})\sim\mathcal{D}} \left[\left(w \log(w) - w \log\left(\min_{j \in [m]} \frac{p_j}{u_j} \right) - w \right) \right], \tag{66}$$

which can be solved to give an optimal price vector $\mathbf{p}^*$. Note that this price vector $\mathbf{p}^*$ corresponds to the static expected equilibrium price, as it takes an expectation over the distribution $\mathcal{D}$. The corresponding pricing policy $\boldsymbol{\pi}$ only depends on the distribution $\mathcal{D}$ is thus given by $\mathbf{p}^t = \mathbf{p}^* = \pi_t(\mathcal{D})$ for all users $t \in$

$[n]$. Given the price vector $\mathbf{p}^*$, all arriving users will purchase an affordable utility-maximizing bundle of goods by solving their individual optimization Problem (1a)-(1c). Note here that the price vector $\mathbf{p}^*$ is computed before the online procedure, which is possible due to the prior knowledge of the distribution $\mathcal{D}$. For numerical implementation purposes, we consider a sample average approximation to compute the expectation in Problem (66), as elucidated in Section 7.2.

Dynamic Learning using SAA: In this benchmark, we consider the setting wherein users' budget and utility parameters are revealed to the central planner each time a user arrives. In this context, the prices are set based on the dual variables of the capacity constraints of the sampled Eisenberg-Gale program with the observed budget and utility parameters of agents that have previously arrived. That is, the pricing policy $\boldsymbol{\pi}$ depends on the history of users' budget and utility parameters, i.e., $\mathbf{p}^t = \pi_t((w_{t'}, \mathbf{u}_{t'})_{t'=1}^{t-1})$. We note that to improve on the computational complexity, we update the dual prices at geometric intervals, as in earlier work (Li and Ye 2022, Agrawal et al. 2014). Users arriving in each interval observe the corresponding price vector for that interval and solve their individual optimization problems to obtain their most favorable goods under the set prices. This process is presented formally in Algorithm 3.

Algorithm 3: Dynamic Learning SAA Algorithm

Input : Vector of Capacities $\mathbf{c}$

Set $\delta \in (1, 2]$ and $L > 0$ such that $\lfloor \delta^L \rfloor = n$;

Let $t_k = \lfloor \delta^k \rfloor, k = 1, 2, \dots, L-1$ and $t_L = n+1$;

Initialize $\mathbf{p}^{t_1} > \mathbf{0}$;

Each user $t \in [t_1]$ purchases a bundle of goods $\mathbf{x}_t$ by solving Problem (1a)-(1c) given the price $\mathbf{p}^{t_1}$;

for $k = 1, \dots, L-1$ **do**

Phase I: Set Price for Geometric Interval

Set price $\mathbf{p}^{t_k}$ based on dual variables of the capacity constraints of the sampled social optimization problem:

$$\begin{aligned} \max_{\mathbf{x}_t \in \mathbb{R}^m, \forall t \in [t_k]} \quad & U(\mathbf{x}_1, ..., \mathbf{x}_{t_k}) = \sum_{t=1}^{t_k} w_t \log \left(\sum_{j=1}^m u_{tj} x_{tj} \right), \\ \text{s.t.} \quad & \sum_{t=1}^{t_k} x_{tj} \leq \frac{t_k}{n} c_j, \quad \forall j \in [m], \\ & x_{tj} \geq 0, \quad \forall t \in [t_k], j \in [m], \end{aligned}$$

Phase II: Each User in Interval Consumes Optimal Bundle

Each user $t \in \{t_k + 1, \dots, t_{k+1}\}$ purchases an optimal bundle of goods $\mathbf{x}_t$ by solving Problem (1a)-(1c) given the price $\mathbf{p}^{t_k}$;

end

Appendix M: Additional Numerical Experiments

M.1. Numerical Validation of Lipschitzness Relation

In this section, we present the results of a numerical experiment to validate the Lipschitzness relation established in Lemma 3. In particular, we consider the instance described in the proof of Theorem 1 with $n = 10,000$ users, where all users have a fixed budget of one, and two goods, each with a capacity of $c_j = n = 10,000$. The utility parameters of users are drawn i.i.d. from a distribution $\mathcal{D}$, where users have an equal 0.5 probability of having the utility (1, 0) or (0, 1).

Figure 2 depicts the change in the dual prices of the certainty equivalent problem between subsequent iterations of Algorithm 1 for this instance. To see that the Lipschitzness relation is satisfied, first note that the norm of the difference between the average remaining resource capacities between subsequent time steps is $O(\frac{1}{n-t})$, i.e., $\|\mathbf{d}_{t+1} - \mathbf{d}_t\|_2 \leq O(\frac{1}{n-t})$ as $\mathbf{d}_{t+1} = \mathbf{d}_t + \frac{\mathbf{d}_t - \mathbf{x}_k(\mathbf{p}_t)}{n-t}$. Then, Figure 2 implies that the two norm of the change in the dual prices of the certainty equivalent problem, i.e., $\mathbb{E}[\|\mathbf{p}^{t+1} - \mathbf{p}^t\|]$ is always upper bounded by $O(\frac{1}{n-t})$ for all $t \in [n-1]$, which thus implies that the obtained Lipschitzness relation in Lemma 3 is satisfied. We note that we present the results on a log plot for readability purposes.

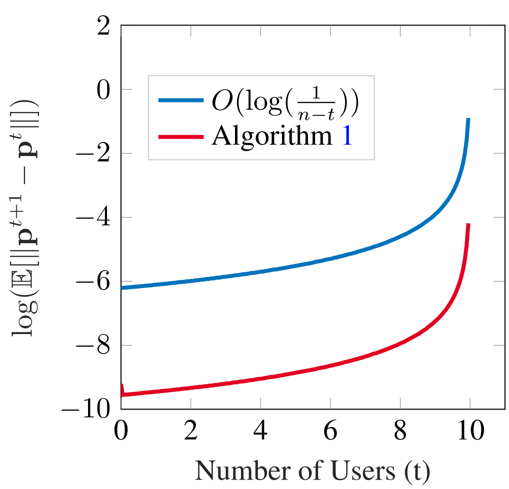

Figure 2 Validation of Lemma 3 for an instance with $n = 10,000$ users, where all users have a fixed budget of one, and two goods, each with a capacity of $c_j = n = 10,000$.

M.2. Numerical Experiments Comparing Static and Adaptive Variants of Expected Equilibrium Pricing

In this section, we numerically evaluate the performance of the static expected equilibrium pricing algorithm and its dynamic counterpart (Algorithm 1) on the counterexample in the proof of Theorem 1. In particular, we considered a setting of n users, where all users have a fixed budget of one, and two goods, each with a capacity of n. The utility parameters of users are drawn i.i.d. from a distribution $\mathcal{D}$, where users have an equal 0.5 probability of having the utility $(1,0)$ or $(0,1)$. For the experiments, we let the number of users n range between 100 to $10,000$ users.

Figure 3 depicts both the constraint violation and the regret of the two algorithms. From the figure, it can be observed that the static expected equilibrium pricing approach achieves negative regret for a large constraint violation, while Algorithm 1 achieves a small positive regret for almost no constraint violation. Recall here from the proof of Theorem 1 that the expected optimal social welfare objective $\mathbb{E}[U_n^*] \in [n\log(2)-1, n\log(2)]$, and thus a regret of less than five for 10,000 users is negligible. As a result, Figure 3 clearly depicts the benefit of adaptivity in online Fisher markets.

We also note that the regret of the static expected equilibrium pricing algorithm is in the range $[-1,0]$, as the accumulated online objective is $n\log(2)$, as each user obtains two units of the good for which they have positive utility under the static expected equilibrium prices of $(0.5,0.5)$ for this instance. As a result, observe that the numerically observed regret in the range $[-1,0]$ aligns with the tight bound for the expected optimal social welfare objective obtained in the proof of Theorem 1, i.e., $n\log(2)-1 \le \mathbb{E}[U_n^*] \le n\log(2)$.

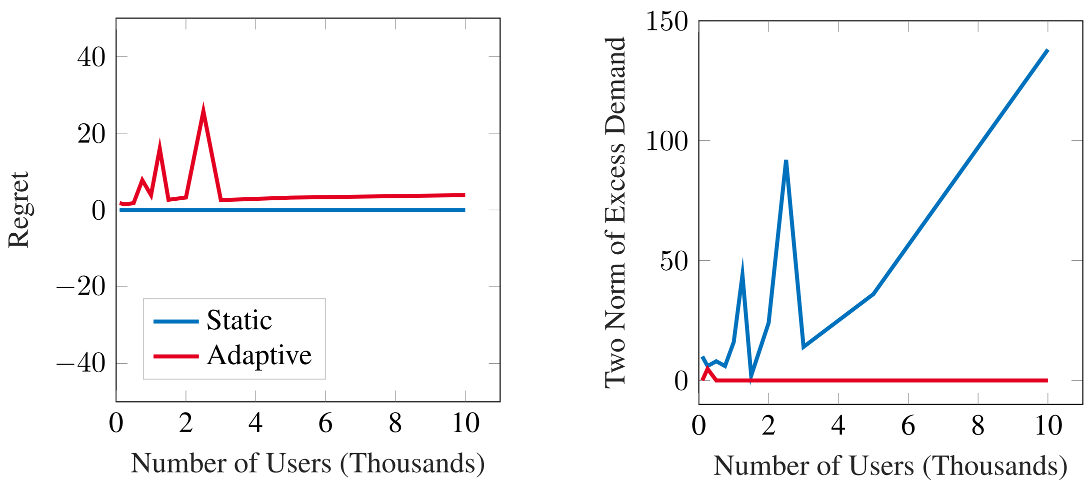

Figure 3 Comparison between the static expected equilibrium pricing algorithm and its dynamic counterpart (Algorithm 1) on regret and constraint violation metrics.

M.3. Numerical Comparison between the Additive and Multiplicative Price Updates in Algorithm 2

We now compare Algorithm 2 (with a fixed step size) that has an additive price update step to a corresponding algorithm with a multiplicative price update step (see Section 5.2). To this end, we consider instance two described in Section 7.2 with a step size of $\gamma = \frac{1}{10\sqrt{n}}$.

Figure 4 depicts the regret and constraint violation for algorithms with the two price update steps given an initial price of 0.5 for all goods. We can observe from Figure 4 that Algorithm 2 with an additive price update rule has a higher regret but a lower constraint violation as compared to the corresponding algorithm with a multiplicative price update rule. This observation highlights the fundamental trade-off between the regret and constraint violation metrics. Yet, we note that since the multiplicative price update rule achieves a lower regret (despite achieving a higher constraint violation) compared to the additive price update rule in Algorithm 2, our results motivate a deeper study of the regret and constraint violation bounds under the multiplicative price update rule.

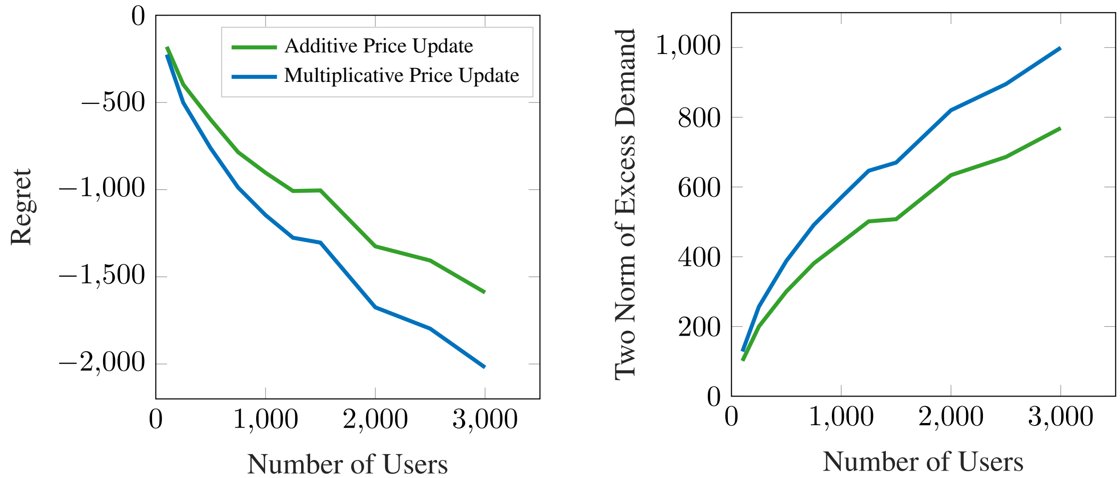

Figure 4 Comparison between Algorithm 2 that has an additive price update step to a corresponding algorithm with a multiplicative price update step on regret and constraint violation metrics.

M.4. Numerical Validation of Positivity of Prices in Algorithm 2

In this section, we present the results of a numerical experiment to validate that the prices remain strictly positive throughout the operation of Algorithm 2 with a fixed step size. To this end, we consider two market settings: (i) the setting described in the counterexample in the proof of Theorem 1, and (ii) instance two described in Section 7.2. For the experiments, we let the number of users n range between 500 to $10,000$ users, consider a step-size of the price updates as $\gamma = \frac{1}{200\sqrt{n}}$, and compute the minimum prices across all goods for 300 instances. In particular, Figure 5 depicts the minimum prices of all goods across 300 instances, which validates the positivity of the prices during the operation of Algorithm 2.

Appendix N: Relation to Approximate Equilibria

In this section, we present the connection between our studied performance metrics and market equilibria. To this end, we first note that our regret and constraint violation metrics approximate the optimal offline Eisenberg Gale aggregated social objective and constraint satisfaction, respectively, which, as beautifully proven, corresponds to perfect Pareto efficiency and envy-freeness under complete information of the utility and budget parameters of users. As a result, obtaining sub-linear guarantees for our regret and constraint violation metrics serves as a proxy for a solution corresponding to an approximate market equilibrium, as the

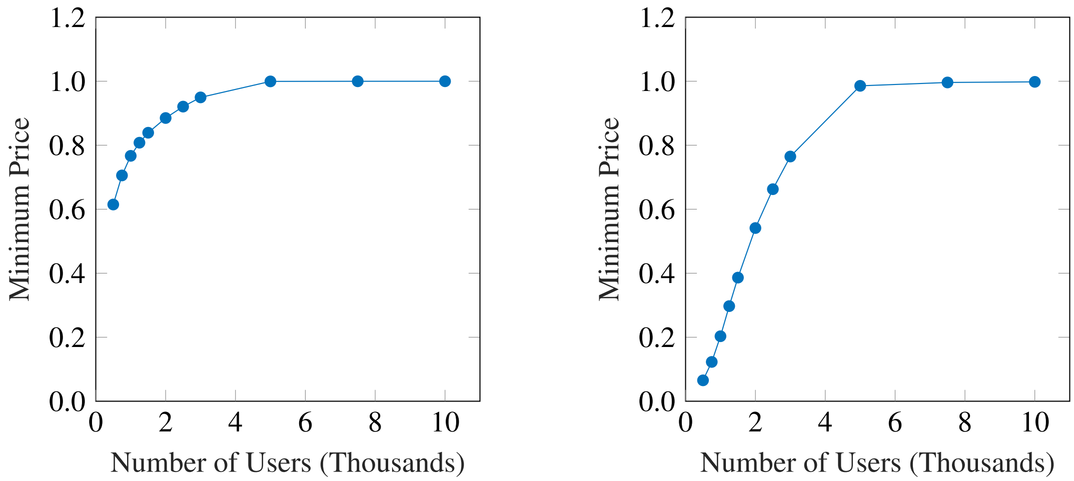

Figure 5 Numerical validation of the positivity of prices during the operation of Algorithm 2 in two market settings: (i) the market instance in the proof of Theorem 1 (left), and (ii) instance two described in Section 7.2 (right). The y-axis denotes the minimum price across all goods across 300 problem instances, i.e., 300 runs of Algorithm 2 on different instances drawn from the specified distribution corresponding to each market setting.

distance to the optimal offline objective and constraint satisfaction of an algorithm indicate its proximity to the optimal offline equilibrium solution. For instance, we note that the per-period regret of Algorithm 2 with a fixed step size of $O(\frac{1}{\sqrt{n}})$ is $\frac{1}{\sqrt{n}}$, which decays and approaches zero as the number of users becomes large. This fact suggests that, on average, the allocations made by the online algorithm approach that of the optimal offline solution as the number of users becomes large, further suggesting that the price-iterates approach the market equilibrium in expectation. Furthermore, we reiterate that achieving low regret corresponding to the Eisenberg Gale objective implies that no user can suffer too much, i.e., receive very low utilities, as the objective is a (weighted) product of all users' utilities.

Our studied problem setting and corresponding performance metrics directly relate to notions of Pareto efficiency and envy-freeness. To this end, we first note that our constraint violation metric can serve as a measure of Pareto inefficiency, which is typically related to the extent to which the capacity constraints are not satisfied (e.g., see Sinclair et al. (2022)), i.e., the number of unsold goods, as when certain goods are unsold some users can become better off without making others worse off. Noting that our theoretical guarantees for constraint violation hold for both the settings of over or under-consumption of resources, our constraint violation bounds thus serve as a measure of the degree of Pareto inefficiency of our obtained solution, which is sub-linear in the number of users for our proposed algorithms.

As for envy-freeness, we first note that our proposed algorithms correspond to posted-price mechanisms, wherein users observe the posted prices and freely (and truthfully) choose which goods to purchase to obtain their most favored bundle of goods given the set prices. In this regard, our proposed algorithms are envy-free by design as all users obtain their most favored bundle of goods, given the set prices upon their arrival. Furthermore, even though under our algorithms, users typically observe different prices, we note that most users observe prices that are similar to other users implying an envy-freeness with regards to the prices faced by users, i.e., there is little that users can gain by swapping their observed prices with that faced by most other users. To elucidate this point, we conducted numerical experiments of Algorithm 2 with a fixed step size under the two market settings: (i) the setting described in the counterexample in the proof of Theorem 1, and (ii) instance two described in Section 7.2. For the experiments, we let the number of users n be $5,000$ and consider a step-size of the price updates as $\gamma = \frac{1}{5\sqrt{n}}$. Figure 6 depicts the evolution of the prices of the goods under both our market instances and demonstrates that a majority of (about 90% of) the users observe prices within a small price band under both market instances. In general, we note that given the $O(\frac{1}{\sqrt{n}})$ step-size of the price updates in Algorithm 2, one can expect that it will take about $O(\sqrt{n})$ steps for the price to move from the initial price vector to a new price vector that, from that point on, stabilizes in a particular band. Thus, Algorithm 2 (with a fixed step size) can be interpreted as achieving approximate envy-freeness where only $O(\sqrt{n})$ users observe arbitrary prices, while the remaining users observe prices within a small price band.

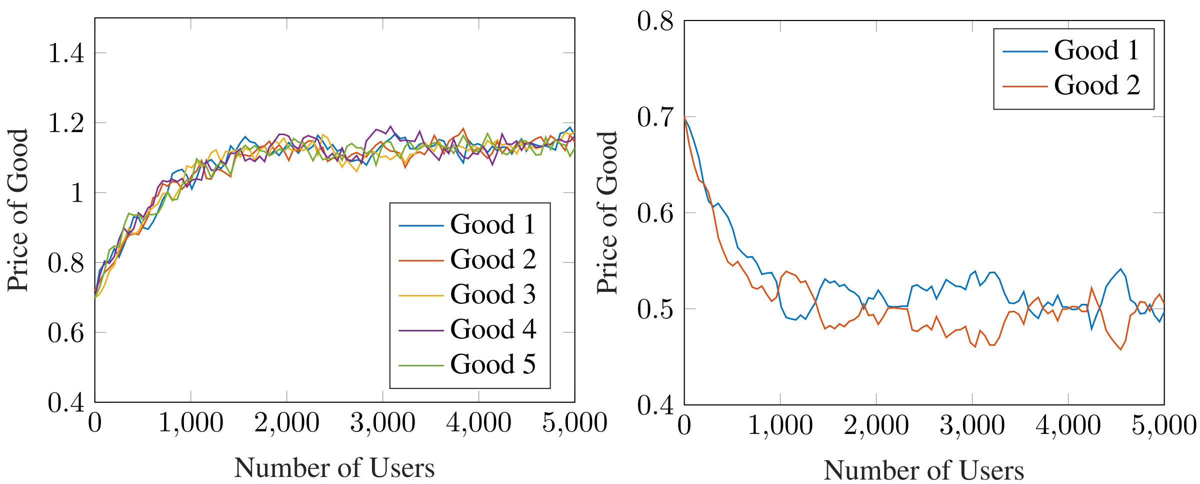

Figure 6 Evolution of the price vector of Algorithm 2 (with a fixed step size) with the number of users on the instance two described in Section 7.2 (left) and the setting described in the counterexample in the proof of Theorem 1 (right). In both market instances, the price vectors stabilize in a small band after the arrival of the first few users.

References of Appendix

[Appendix3] Agrawal S, Wang Z, Ye Y (2014) A dynamic near-optimal algorithm for online linear programming. *Operations Research* 62(4):876–890.

[Appendix2] Balcan MF, Daniely A, Mehta R, Urner R, Vazirani V (2014) Learning economic parameters from revealed preferences. Liu TY, Qi Q, Ye Y, eds., *International Conference on Web and Internet Economics*, 338–353 (Cham: Springer International Publishing), ISBN 978-3-319-13129-0.

[Appendix3] Bei X, Chen W, Garg J, Hoefer M, Sun X (2016) Learning market parameters using aggregate demand queries. *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 404–410, AAAI'16 (Phoenix, Arizona: AAAI Press).

[Appendix4] Beigman E, Vohra R (2006) Learning from revealed preference. *Proceedings of the 7th ACM Conference on Electronic Commerce*, 36–42, EC '06, ISBN 1595932364.

[Appendix5] Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103.

[Appendix6] Cao X, Liu KJR (2019) Online convex optimization with time-varying constraints and bandit feedback. *IEEE Transactions on Automatic Control* 64(7):2665–2680.

[Appendix12] Chen G, Li X, Ye Y (2024) Technical note—an improved analysis of lp-based control for revenue management. *Operations Research* 72(3):1124–1138, URL http://dx.doi.org/10.1287/opre.2022.2358.

[Appendix8] Chen T, Ling Q, Giannakis GB (2017) An online convex optimization approach to proactive network resource allocation. *IEEE Transactions on Signal Processing* 65(24):6350–6364.

[Appendix9] Devanur N, Hayes T (2009) The adwords problem: Online keyword matching with budgeted bidders under random permutations. *Proceedings of the 10th ACM Conference on Electronic Commerce*, 71–78, EC '09, ISBN 9781605584584.

[Appendix18] Gao W, Sun C, Xue C, Ge D, Ye Y (2024) Decoupling learning and decision-making: Breaking the $\mathcal{O}(\sqrt{T})$ barrier in online resource allocation with first-order methods.

[Appendix11] Gorokh A, Banerjee S, Iyer K (2016) Near-efficient allocation using artificial currency in repeated settings. *Available at SSRN 2852895* .

[Appendix12] Gorokh A, Banerjee S, Iyer K (2021) From monetary to nonmonetary mechanism design via artificial currencies. *Mathematics of Operations Research* 46(3):835–855.

[Appendix13] Jalota D, Pavone M, Qi Q, Ye Y (2020) Markets for efficient public good allocation with social distancing. *Web and Internet Economics: 16th International Conference, WINE 2020, Beijing, China, December 7–11, 2020, Proceedings*, 102–116 (Berlin, Heidelberg: Springer-Verlag), ISBN 978-3-030-64945-6, URL http://dx.doi.org/10.1007/978-3-030-64946-3_8.

[Appendix25] Jenatton R, Huang J, Archambeau C (2016) Adaptive algorithms for online convex optimization with long-term constraints. *International Conference on Machine Learning*, 402–411 (PMLR).

[Appendix26] Ji Z, Mehta R, Telgarsky M (2018) Social welfare and profit maximization from revealed preferences. Christodoulou G, Harks T, eds., *Web and Internet Economics*, 264–281 (Cham), ISBN 978-3-030-04612-5.

[Appendix16] Kash IA, Friedman EJ, Halpern JY (2007) Optimizing scrip systems: Efficiency, crashes, hoarders, and altruists. *Proceedings of the 8th ACM Conference on Electronic Commerce*, 305–315, EC '07 (New York, NY, USA: Association for Computing Machinery), ISBN 9781595936530, URL http://dx.doi.org/10.1145/1250910.1250955.

[Appendix28] Kleinberg R, Leighton T (2003) The value of knowing a demand curve: bounds on regret for online posted-price auctions. *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, 594–605.

[Appendix18] Lalley SP (2013) Concentration inequalities. *Lecture notes, University of Chicago* .

[Appendix30] Li X, Sun C, Ye Y (2022) Simple and fast algorithm for binary integer and online linear programming. *Math. Program.* 200(2):831–875, ISSN 0025-5610, URL http://dx.doi.org/10.1007/s10107-022-01880-x.

[Appendix31] Li X, Ye Y (2022) Online linear programming: Dual convergence, new algorithms, and regret bounds. *Operations Research* 70(5):2948–2966, URL http://dx.doi.org/10.1287/opre.2021.2164.

[Appendix21] Liakopoulos N, Destounis A, Paschos G, Spyropoulos T, Mertikopoulos P (2019) Cautious regret minimization: Online optimization with long-term budget constraints. *International Conference on Machine Learning*, 3944–3952 (PMLR).

[Appendix32] Mahdavi M, Jin R, Yang T (2012) Trading regret for efficiency: online convex optimization with long term constraints. *The Journal of Machine Learning Research* 13(1):2503–2528.

[Appendix23] Prendergast C (2016) The allocation of food to food banks. *EAI Endorsed Trans. Serious Games* 3(10):e4.

[Appendix36] Roth A, Ullman J, Wu ZS (2016) Watch and learn: Optimizing from revealed preferences feedback. *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, 949–962, STOC '16, ISBN 9781450341325.

[Appendix37] Sinclair S, Jain G, Banerjee S, Yu CL (2022) Sequential fair allocation: Achieving the optimal envy-efficiency trade-off curve. *Operations Research* 0(0):null.

[Appendix26] Valls V, Iosifidis G, Leith D, Tassiulas L (2020) Online convex optimization with perturbed constraints: Optimal rates against stronger benchmarks. *International Conference on Artificial Intelligence and Statistics*, 2885–2895 (PMLR).

[Appendix27] van de Geer S (2002) *On Hoeffding's Inequality for Dependent Random Variables*, 161–169 (Boston, MA: Birkhäuser Boston), ISBN 978-1-4612-0099-4.

[Appendix28] Yi X, Li X, Yang T, Xie L, Chai T, Johansson K (2021) Regret and cumulative constraint violation analysis for online convex optimization with long term constraints. *International Conference on Machine Learning*, 11998–12008 (PMLR).

[Appendix45] Yu H, Neely M, Wei X (2017) Online convex optimization with stochastic constraints. *Advances in Neural Information Processing Systems* 30.

[Appendix30] Zadimoghaddam M, Roth A (2012) Efficiently learning from revealed preference. *Proceedings of the 8th International Conference on Internet and Network Economics*, 114–127, WINE'12 (Berlin, Heidelberg: Springer-Verlag), ISBN 9783642353109.