Online Market Intermediation*†

Yiannis Giannakopoulos[1], Elias Koutsoupias[2], and Philip Lazos[3]

1 **Department of Informatics, TU Munich, Munich, Germany**
giannako@in.tum.de
2 **Department of Computer Science, University of Oxford, Oxford, UK**
elias@cs.ox.ac.uk
3 **Department of Computer Science, University of Oxford, Oxford, UK**
filippos.lazos@cs.ox.ac.uk

Abstract

We study a dynamic market setting where an intermediary interacts with an unknown large sequence of agents that can be either sellers or buyers: their identities, as well as the sequence length n, are decided in an adversarial, online way. Each agent is interested in trading a single item, and all items in the market are identical. The intermediary has some prior, incomplete knowledge of the agents' values for the items: all seller values are independently drawn from the same distribution F_S, and all buyer values from F_B. The two distributions may differ, and we make common regularity assumptions, namely that F_B is MHR and F_S is log-concave.

We focus on online, posted-price mechanisms, and analyse two objectives: that of maximizing the intermediary's profit and that of maximizing the social welfare, under a competitive analysis benchmark. First, on the negative side, for general agent sequences we prove tight competitive ratios of $\Theta(\sqrt{n})$ and $\Theta(\ln n)$, respectively for the two objectives. On the other hand, under the extra assumption that the intermediary knows some bound α on the ratio between the number of sellers and buyers, we design asymptotically optimal online mechanisms with competitive ratios of $1 + o(1)$ and 4, respectively. Additionally, we study the model where the number of items that can be stored in stock throughout the execution is bounded, in which case the competitive ratio for the profit is improved to $O(\ln n)$.

1998 ACM Subject Classification J.4 [Social and Behavioral Sciences] Economics

Keywords and phrases optimal auctions, bilateral trade, sequential auctions, online algorithms, competitive analysis

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.47

1 Introduction

The design and analysis of electronic markets is of central importance in algorithmic game theory. Of particular interest are trading settings, where multiple parties such as buyers, sellers, and intermediaries exchange goods and money. Typical examples are markets for trading stocks, commodities, and derivatives: sellers and buyers where each one trades a single item and one intermediary for facilitating the transactions. However, the well-understood cases are comparatively quite modest. The very special case of one seller, and one buyer was thoroughly studied by Myerson and Satterthwaite [26] in their seminal paper; they provided a beautiful characterization of many significant properties a mechanism might have, along

* A full version of this paper can be found in [18], http://arxiv.org/abs/1703.09279.

† This work was kindly supported by the ERC Advanced Grant 321171 (ALGAME), ERC Advanced Grant 691672 (APEG) and by the EPSRC (award reference 1493310).

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).
Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;
Article No. 47; pp. 47:1–47:14
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

with an impossibility theorem showing that it cannot possess them all. The paper also dealt with the case where a broker provides assistance by making two potential trades, one with each agent, while also trying to maximize his profit. This was extended in [15] to multiple sellers and buyers that are all immediately present in an offline manner.

Our work considers a similar setting, but with a key difference: the buyers and sellers appear one-by-one, in a dynamic way. It is natural to study this question in the incomplete information setting in which the intermediary, whose objective is to maximize either profit or welfare, does not know the sequence of buyers and sellers in advance. The framework that we employ to study the question is the standard worst-case analysis of online algorithms whose goal is to do as well as possible in the face of a powerful adversary which tries to embarrass them.

We are not the first to apply techniques from online algorithms to quantify uncertainty in markets: the closest work to ours would be by Blum et al. [8] who consider buyers and sellers trading identical items. In their setting, motivated mostly from a financial standpoint, buyers and sellers arrived in an online manner, with their bids appearing to the trader and expiring after some time. The trader would have to match prospective buyers and sellers to facilitate trade. Among a plethora of interesting results, the trader's profit maximization problem was studied using competitive analysis and techniques from online weighted matchings. The key difference in our setting is that buyers and sellers do not overlap: whenever a seller appears, the intermediary has to decide whether or not to attempt to buy the item, without having a buyer ready to go. Instead, the intermediary stores the item to sell it at a later time. We believe this variation is able to capture "slower" markets, like online marketplaces similar to Amazon or AliExpress (or even regular retail stores), where uncertainty stems from not knowing how large a stock of items to buy, in expectation of the buyers to come.

1.1 Our Results

Our aim is to study this dynamic market setting, where an intermediary faces a sequence of potential buyers and sellers in an online fashion. The goal of the intermediary is to maximize his profit, or society's welfare, by buying from the sellers and selling to buyers. We take a Bayesian approach to their utilities but use competitive analysis for their arrivals: the main difficulty stems from the unknown (and adversarially chosen) sequence of agents. Further particulars and notation is discussed in Section 2. All the online algorithms we design are posted price, which are simple, robust and strongly truthful.

First, in Section 4 we study the case of arbitrary sequences of buyers and sellers and show that the competitive ratio—the ratio of the optimal offline profit over the profit obtained by the online algorithm—is $\Theta(\sqrt{n})$, where n is the total number of buyers and sellers. We also study the social welfare objective, where the goal is to maximize the total utility of all participants, including the sellers, the buyers and the intermediary The competitive ratio here is $\Theta(\log n)$. All these results are achieved via common regularity assumptions on the distributions of the agent values (see Section 3), which we also prove to be necessary, by providing arbitrarily bad competitive ratios in the case they are dropped (Theorem 7).

To overcome the above pessimistic results, we next study in Section 5 the setting where both the online and offline algorithms have a limited stock, i.e. at no point in time can they hold more than K items. In this model, the competitive ratio is improved to $\Theta(K \log n)$, asymptotically matching that of welfare. Finally, we also propose a way to restrict the input sequence, by introducing in Section 6 the notion of α-balanced streams, where at every prefix of the stream the ratio of the number of sellers to buyers has to be at least α. Under this condition we are able to bring down the competitive ratios for both objectives to constants. In particular, the online posted-price mechanism that we use for profit maximization, and which is derived by a fractional relaxation of the optimal offline profit, achieves an asymptotically optimal ratio of $1 + o(1)$. A similar mechanism is 4-competitive for the welfare objective.

All omitted proofs can be found in the full version of the paper [18].

1.2 Prior Work

Our work is grounded on a string of fruitful research in mechanism design. The main topics that are close to our effort are bilateral trading, trading markets and sequential (online) auctions.

The first step in bilateral trading and mechanism design was made by Myerson and Satterthwaite [26] who proved their famous impossibility result, even for the case of one buyer and one seller. The case for profit maximization was extended to many buyers and sellers, each trading a single identical item, in [15]. Some of the assumptions in our model are based in these two works. The impossibility result in [26], among other difficulties, slowly vanishes for larger markets as was shown by McAffee [25]. There is still active progress being made on this intriguing setting, concentrating on simple mechanisms that provide good approximations either to welfare while staying budget balanced and individually rational [9, 11] or to profit [27]. Other recent developments include a hardness result for computing optimal prices [17] and constant efficiency approximation with strong budget balance [14].

Sequential auctions have also produced a collection of interesting results, either extending the ideas of simple approximate mechanisms instead of more complex, theoretically optimal ones or dealing with entirely new settings. Prominent examples that compare the revenue (or welfare) generated by simple, posted-price sequential auctions to the optimal, proving good approximations in certain cases, are [10] for single-item revenue, [13, 29] for matroid constraints (and some multi-dimensional settings) and [16] for combinatorial auctions. There have been many approaches that apply competitive (worst-case) analysis to mechanism design. The analysis of auctions with unlimited supply is explored in [5, 7] where near optimal algorithms are developed using techniques inspired from no-regret learning. There is also a deep connection between secretary problems and online sequential auctions [21, 20, 3]. Hajiaghayi et al. utilized techniques such as prophet inequalities for unknown market size with distributional assumptions in [22]. A comprehensive exposition of online mechanism design by Parkes can be found in [28].

There are also positive results in online auctions when the valuation distribution is unknown (but usually known to be restricted in some way, having bounded support or being monotone hazard-rate etc). Babaioff et al. explored the case of selling a single item to multiple i.i.d. buyers in [1]. The case of k items in a similar setting was studied in [2], while the case of unlimited items (digital goods auctions) in [23] and [24]. Budget constraints where also introduced in [4], where a procurement auction was the focus.

2 Preliminaries and Notation

The input is a finite string $\sigma \in \{S, B\}^*$ of buyers (B) and sellers (S). The online algorithm has no knowledge of $\sigma(t)$, i.e. whether $\sigma(t) = S$ or $\sigma(t) = B$, before step t. Also, it doesn't know the length $n(\sigma)$ of σ. Denote $n_S(\sigma)$, $n_B(\sigma)$ the number of sellers and buyers, respectively, in σ, and let $N_S(\sigma)$, $N_B(\sigma)$ be the corresponding set of indices, i.e. $N_S(\sigma) = \{t \mid \sigma(t) = S\}$ and $N_B(\sigma) = \{t \mid \sigma(t) = B\}$. Let $N(\sigma) = N_S(\sigma) \cup N_B(\sigma) = \{1, 2, \ldots, n(\sigma)\}$. In the above notation we will often drop the σ if it is clear which input stream we are referring to.

The values of the sellers are drawn i.i.d. from a probability distribution (with cdf) F_S and these of buyers i.i.d. from a distribution F_B, both supported over intervals of nonnegative reals. We denote the random variable of the value of the t-th agent with X_t. We assume that distributions F_S and F_B are continuous, with bounded expectations μ_S and μ_B, and have (well-defined) density functions f_S and f_B, respectively. It will also be useful to denote by X_S a random variable drawn from distribution F_S, and similarly $X_B \sim F_B$, and for any random variable Y and positive integer m use $Y^{(m)}$ to represent the maximum order statistic out of m i.i.d. draws from the same distribution as Y. We will also use the shortcut notation $\mu^{(m)} = \mathbb{E}[Y^{(m)}]$.

We study *posted-price* online algorithms that upon seeing the identity of the t-th agent (whether she is a seller or a buyer), offer a price p_t. We buy one unit of the item from sellers that accept our price (i.e. if $\sigma(t) = S$ and $X_t \le p_t$) and pay them that price, and we sell to buyers that accept our price (i.e. if $\sigma(t) = B$ and $X_t \ge p_t$), given stock availability (see below), and collect from them that price. So, a price p_{t+1} can only depend on $\sigma(1), \dots, \sigma(t+1)$ and the result of the comparison $X_i \le p_i$ in all previous steps $i = 1, 2, \dots, t$. Let K_t denote the available stock at the beginning of the t-th step, i.e. $K_1 = 0$ and

$$K_{t+1} = \begin{cases} K_t + 1, & \text{if } \sigma(t) = S \ \wedge \ X_t \le p_t \\ K_t - 1, & \text{if } \sigma(t) = B \ \wedge \ K_t \ne 0 \ \wedge \ X_t \ge p_t \\ K_t, & \text{otherwise.} \end{cases}$$

Then, the set of sellers from whom we bought items during the algorithm's execution is $I_S = \{t \in N_S \mid X_t \le p_t\}$ and the set of buyers we sold to is $I_B = \{t \in N_B \mid X_t \ge p_t \wedge K_t \ne 0\}$. Notice that these are random variables, depending on the actual realizations of the agent values X_t.

The total *profit* that the intermediary deploying an algorithm A makes throughout the execution on an input stream σ, is the amount he manages to collect from the buyers via successful sales, minus the amount he spent in order to maintain stock availability from the sellers, that is

$$\mathcal{R}(A, \sigma) = \mathbb{E}\left[\sum_{t \in I_B} p_t - \sum_{t \in I_S} p_t\right].$$

The social *welfare* of algorithm A is the sum of valuations that all participants achieve throughout the entire execution. That is, a seller at position t of the stream has a value of X_t if she keeps her item, or a value of p_t if she sold the item to the intermediary; a buyer has a value of $X_t - p_t$ if she managed to buy an item, since the item has a value of X_t and he spent p_t to buy it, or 0 otherwise. And the intermediary, has a value of $\mathcal{R}(A)$ plus the value of the items that he didn't manage to sell in the end and which are now left in his stock. Putting everything together and performing the occurring cancellations, this results in the welfare to be expressed simply as the sum of the values of the sellers that kept their items plus the sum of the values of the buyers that bought an item, i.e.

$$\mathcal{W}(A, \sigma) = \mathbb{E}\left[\sum_{t \in N_S \setminus I_S} X_t + \sum_{t \in I_B} X_t\right]. \tag{1}$$

We use *competitive analysis*, the standard benchmark for online algorithms (see e.g. [12]), in order to quantify the performance of an online algorithm A: we compare it to that of an unrealistic, offline optimal algorithm OPT has access to the entire stream σ in advance.

Then, we say that A is $\rho(n)$-competitive with respect to welfare, if for any feasible input sequence of agents σ with length n and distributions F_S, F_B for the agent values, it is $\mathcal{W}(\mathrm{OPT},\sigma) \leq \rho(n) \cdot \mathcal{W}(A,\sigma)$. Notice how we allow the competitive ratio $\rho(n)$ to explicitly depend on the input's length, so that we can perform asymptotic analysis as $\mathcal{W}(\mathrm{OPT},\sigma)$ and n tend to infinity. It is common in competitive analysis to allow for an additional constant in the right hand side of the above expression, that does not depend in the input, and which intuitively can capture some initial configuration disadvantage of the online algorithm. We do that for the case of the profit objective, as this constant will have a very natural interpretation: you can think of it as the maximum amount of deficit on which an online algorithm can run at any point in time, since an adversary can always stop the execution at any time he wishes. Given that interpretation, it makes sense to allow for this constant to depend on seller distribution F_S, since even when we face a single seller at the first step we expect to spend an amount that depends on the realization of her value. Thus, we will say that an online algorithm is $\rho(n)$-competitive with respect to profit, if for any input sequence of agents σ and any probability priors F_S, F_B,

$$\mathcal{R}(\mathrm{OPT},\sigma) \leq \rho(n) \cdot \mathcal{R}(A,\sigma) + O(\mu_S). \tag{2}$$

3 Distributional Assumptions

Throughout most of the paper we will make some assumptions on the distributions F_B, F_S from which the buyer and seller values are drawn. In particular, we will assume that F_B has *monotone hazard rate (MHR)*, i.e. $\log(1 - F_B(x))$ is concave, and that F_S is *log-concave*, i.e. $\log F_S(x)$ is concave. For convenience, we will collectively refer to both the above constraints as *regularity assumptions*. These conditions are rather standard in the optimal auctions literature, and they encompass a large class of natural of distributions including e.g. exponential, uniform and normal ones. Notice that distributions that satisfy the above conditions also fulfil the regularity requirements introduced in the seminal paper Myerson and Satterthwaite [26] for the single-shot, one buyer and one seller setting of bilateral trade, namely that $x + \frac{F_S(x)}{f_S(x)}$ and $x - \frac{1-F_B(x)}{f_B(x)}$ are both increasing functions. Finally, we must mention that such regularity assumptions are necessary, in the sense that dropping them would result in arbitrarily bad lower bounds for the competitive ratios of our objectives, as it is demonstrated by Theorem 7.

The following two lemmas demonstrate some key properties of distributions satisfying our regularity assumptions and which will be very useful in our subsequent analysis:

▶ **Theorem 1.** *For any random variable Y drawn from an MHR distribution with bounded expectation μ and standard deviation s,*

1. $\Pr[Y \geq y] \geq \frac{1}{e}$ *for any* $y \leq \mu$
2. $\Pr[Y \geq y] < \frac{1}{e}$ *for any* $y > 2\mu$
3. $\mathbb{E}[Y^{(m)}] \leq H_m \cdot \mu$*, where* H_m *is the m-th harmonic number.*
4. $s \leq \mu$

Proof. A proof of Property 1 can be found in [6, Theorem 3.8], of Property 2 in [6, Corollary 3.10], and of Property 3 in [1, Lemma 13]. For Property 4, from [19, Lemma 2] we know that $\mathbb{E}\left[Y^2\right] \leq 2\mu^2$, so $s^2 = \mathbb{E}\left[Y^2\right] - \mu^2 \leq \mu^2$. ◀

▶ **Lemma 2.** *For any distribution over $[0,\infty)$ with log-concave cdf F and expectation μ,*

$$x \leq e\mu F(x) \qquad \text{for any } x \leq \mu.$$

Finally, we prove the following property bounding the sum of maximum order statistics of a distribution, that holds for general (not necessarily MHR) distributions and might be of independent interest:

▶ **Lemma 3.** *The expected average of the k-th highest out of m independent draws from a probability distribution with expectation μ and standard deviation s can be at most $\mu + 2\sqrt{\frac{m}{k}}s$.*

4 General Setting

We start by studying the general setting where no additional assumptions are enforced on the structure of the input sequence. The adversary is free to arbitrarily choose the identities of the agents.

4.1 Welfare

▶ **Theorem 4.** *Under our regularity assumptions*[1]*, the online auction that posts to any seller and buyer the median of their distribution is $O(\ln n)$-competitive with respect to welfare. This bound is tight.*

Proof. We split the proof of the theorem in two more general lemmas below, corresponding to upper and lower bounds. Then, the upper bound for our case follows easily from Lemma 5 by using constants $c_1 = c_2 = 2$, and taking into consideration that, from Property 3 of Theorem 1, the ratio of the maximum order statistic for the MHR distribution F_B is upper bounded by $r_B(m) \leq H_m \leq O(\ln m)$. For the lower bound, it is enough to observe that this ratio is attained by an exponential distribution, which is MHR.

▶ **Lemma 5.** *For any choice of constants $c_1, c_2 > 1$, the following fixed-price online auction has a competitive ratio of at most $\max\left\{\frac{c_1}{c_1-1}, c_1c_2 \cdot r_B(n_B)\right\}$ with respect to welfare, where n_B is the number of buyers, and $r_B(m) = \mu_B^{(m)}/\mu_B$ is the ratio between the m-maximum-order statistic and the expectation of the buyer value distribution.*

- *Post to all sellers price $q = F_S^{-1}\left(\frac{1}{c_1}\right)$.*
- *Post to all buyers price $p = F_B^{-1}\left(\frac{c_2-1}{c_2}\right)$.*

Proof. Let A denote our online algorithm and OPT an offline algorithm with optimal expected welfare. Fix an input stream σ. Looking at (1), the maximum welfare that OPT can get from the sellers is at most $\mathbb{E}\left[\sum_{t\in N_S} X_t\right] = n_s\mu_S$, while from the buyers at most $\mathbb{E}\left[|I_B| \cdot X_B^{(n_B)}\right] \leq \kappa\,\mathbb{E}\left[X_B^{(n_B)}\right]$, where κ is the maximum number of sellers that can be matched to *distinct* buyers that arrive after them[2] in σ: clearly, no mechanism can sell more than κ items. Bringing all together we have that

$$\mathcal{W}(\text{OPT}) \leq n_s\mu_S + \kappa\mu_B^{(n_B)} = n_s\mu_S + r_B(n_B) \cdot \kappa\mu_B.$$

---

[1] As matter of fact, in the proof of Theorem 4 just regularity for the buyer values would suffice, i.e. F_B being MHR.

[2] You can think of that as the maximum size of a matching in the following undirected graph: the nodes are the sellers and the buyers, and there is an edge between any seller and all the buyers that appear after her in σ.

For the online algorithm now, from the sellers we get

$$\sum_{i \in N_S} \Pr\left[X_i > q\right] \mathbb{E}[X_i | X_i > q] \geq n_s(1 - F_S(q)) \, \mathbb{E}[X_S] = \frac{c_1 - 1}{c_1} \cdot n_S \mu_S$$

and from the buyers at least

$$\kappa \Pr\left[X_S \leq q\right] \Pr\left[X_B \geq p\right] \mathbb{E}[X_i | X_i \geq p] \geq \kappa F_S(q)(1 - F_B(p)) \, \mathbb{E}[X_B] = \frac{1}{c_1}\frac{1}{c_2} \cdot \kappa \mu_B,$$

just by considering one of the κ-size matchings discussed before: if we manage to buy from one of these κ sellers, then we will definitely have stock availability for the matched buyer. ◀

The upper bound in Lemma 5 cannot be improved:

▶ **Lemma 6.** *For* any *probability distribution F, even if the seller and buyer values are i.i.d. from F, the sequence SB^n forces all posted-price online mechanisms to have a competitive ratio of $\Omega(r(n))$, where $r(n) = \mu^{(n)}/\mu$ is the ratio of the n-maximum-order statistic of distribution F to its expectation.* ◀

As the following theorem demonstrates, our regularity assumption on the agent values is necessary if we want to hope for non-trivial bounds. In particular, the lower bound in Lemma 6 can be made arbitrarily high:

▶ **Theorem 7.** *For any constant $\varepsilon \in (0, 1)$, there exists a continuous probability distribution F such that any online posted-price mechanism has a competitive ratio of $\Omega(n^{1-\varepsilon})$ on the input sequence SB^n, even if the values of the sellers and the buyers are i.i.d.*

4.2 Profit

Now we turn our attention to our other objective of interest, that of maximizing the expected profit of the intermediary. As it turns out, this objective has some additional challenges that we need to address. For example, as the following theorem demonstrates, if the distribution of seller values is bounded away from 0, the competitive ratio can be arbitrarily bad, even for i.i.d. values from a uniform distribution. Intuitively, this follows from the impossibility of buying a super-constant number of items within a constant budget.

▶ **Theorem 8.** *For any $a > 0$ and $\varepsilon \in (0, 1)$, if the seller and buyer values are drawn i.i.d. from the uniform distribution over $[a, b]$ where $b > 2a$, then no online posted-price mechanism can have an approximation ratio better than a $\left(1 - \frac{1}{k}\right)^4 n^{1-\varepsilon}$ with respect to profit, where $k = \frac{b}{a} - 1$. In particular, for any uniform distribution over an interval $[1, h]$ with $h \geq 3$ the lower bound is $\frac{1}{2^4} n^{1-\varepsilon} = \Omega\left(n^{1-\varepsilon}\right)$.*

If we consider distributions supported over intervals that include 0, under our regularity assumptions we can do a little better than the trivial lower bound of Theorem 8:

▶ **Theorem 9.** *Under our regularity assumptions, for agent values distributed over intervals that include 0 the following online posted-price mechanism achieves a competitive ratio of $O(n^{\frac{1}{2}+\varepsilon})$ for any $\varepsilon > 0$:*

- *Post to the i-th seller price $q_i = F_S^{-1}\left(\frac{1}{e}\frac{1}{i^{1/2+\varepsilon}}\right)$*
- *Post to all buyers price $p = \mu_B$.*

Proof. Fix an input stream σ of length n. Let μ_B and s_B be the expectation and standard deviation of the buyer value distribution F_B. As in the proof of Lemma 5, let κ denote the maximum number of sellers that can be matched to distinct buyers that arrive after them in σ. If $\mu_B^{(j:m)}$ denotes the expectation of the j-th largest out of m independent draws from F_B, since no algorithm can make more than κ sales over its entire execution, the optimal offline profit is upper bounded by

$$\sum_{j=1}^{\kappa} \mu_B^{(n_B-j+1:n_B)} \leq \sum_{i=n-\kappa+1}^{n} \mu_B^{(i:n)} \leq \kappa\mu_B + 2\sqrt{\kappa n} s_B \leq 3\sqrt{\kappa}\sqrt{n}\mu_B,$$

where for the second inequality we have used Lemma 3 and for the last one we have used Property 4 from Theorem 1 and the obvious fact that $\kappa \leq n$.

For the analysis of the online mechanism now, the expected number of items that it gets from the first κ sellers is $\sum_{i=1}^{\kappa} F_S(q_i) = \frac{1}{e}\sum_{i=1}^{\kappa} \frac{1}{i^{1/2+\varepsilon}} \geq \frac{1}{e}\kappa^{1/2-\varepsilon}$. So, by considering the FIFO matching between these first κ sellers and their corresponding buyers, the expected income of our algorithm is at least $\frac{1}{e}\kappa^{1/2-\varepsilon}(1-F(p)) = \frac{1}{e}\kappa^{1/2-\varepsilon}(1-F(\mu_B)) \geq \frac{1}{e^2}\kappa^{1/2-\varepsilon}$, where in the last step we deployed Property 1 of Theorem 1. So, it only remains to be shown that the online algorithm does not spend more than a constant amount. Indeed, our expected spending is at most

$$\sum_{i=1}^{\infty} q_i F_S(q_i) \leq \sum_{i=1}^{\infty} e\mu_S F_S(q_i)^2 = \frac{1}{e}\mu_S \sum_{i=1}^{\infty} \frac{1}{i^{1+2\varepsilon}} = O(\mu_S),$$

where for the first inequality we have used Lemma 2, taking into consideration that seller prices q_i are decreasing and q_1 is below μ_S. This is true because again from Lemma 2 for $x = \mu_S$ we know that $\mu_S \leq e\mu_S F(\mu_S)$, or equivalently $F(\mu_S) \geq \frac{1}{e} = F(q_1)$. ◀

The algorithm of Theorem 9 is asymptotically optimal:

▶ **Theorem 10.** *If the seller and buyer values are drawn i.i.d. from the uniform distribution over* $[0, 1]$*, then no online posted-price mechanism can have an approximation ratio better than* $\Omega(\sqrt{n})$.

Proof. We use the input sequence $\sigma = S^{n/2}B^{n/2}$ with n even. Let $F(x) = x$ be the cdf of the uniform distribution over $[0, 1]$. This time we argue that no online algorithm can buy more than $\Omega(\sqrt{n})$ items from the sellers, in expectation. Indeed, let q_i be the price that the online mechanism posts to the i-th seller. Then, the expected number of items m_σ bought from the sellers is $\sum_{i=1}^{n/2} F(q_i) = \sum_{i=1}^{n/2} q_i$, while the expected expenditure c_σ is $\sum_{i=1}^{n/2} F(q_i)q_i = \sum_{i=1}^{n/2} q_i^2$. By the convexity of the function $t \mapsto t^2$ and Jensen's inequality it must be that

$$m_\sigma = \sum_{i=1}^{n/2} q_i \leq \sqrt{\frac{n}{2}} \left(\sum_{i=1}^{n/2} q_i^2\right)^{\frac{1}{2}} = O\left(\sqrt{c_\sigma}\sqrt{n}\right),$$

so given that our deficit must be $c_\sigma = O(\frac{1}{2})$, we get the desired $m_\sigma = O(\sqrt{n})$. As a result, the online profit can be at most $O(\sqrt{n}) \cdot 1 = O(\sqrt{n})$.

For the offline algorithm we use prices $q = \frac{1}{8}$ and $p = \frac{1}{2}$ for the buyers and sellers, respectively, and by an analogous analysis to that of the proof of Theorem 8, we get that the expected offline profit is at least

$$\frac{n}{2}F(q)(1-F(p))p - \frac{n}{2}F(q)q = \frac{n}{2}\frac{1}{8}\left(1-\frac{1}{2}\right)\frac{1}{2} - \frac{n}{2}\frac{1}{8}\frac{1}{8} = \frac{n}{128} = \Omega(n).$$ ◀

5 Limited Stock

If one looks carefully at the lower bound proof for the profit in Theorem 10, it becomes clear that the source of difficulty for any online algorithm is essentially the fact that without knowledge of the future, you cannot afford to spend a super-constant amount of money into accumulating a large stock of items, without the guarantee that there will be enough demand from future buyers. In particular, it may seem that the offline algorithm has an unrealistic advantage of using a stock of infinite size. The natural way to mitigate this would be to introduce an upper bound K on the number of items that both the online and offline algorithms can store at any point in time. As it turns out, this has a dramatic improvement in the competitive ratio for the profit:

▶ **Theorem 11.** *Assuming stock sizes of at most K items, under our regularity assumptions the following online mechanism is $O(Kr\log n)$-competitive, where $r = \max\left\{1, \frac{\mu_S}{\mu_B}\right\}$:*

- *If your stock is not currently full, post to sellers price $q = F_S^{-1}\left(\frac{1}{r}\frac{1}{2eK}\right)$*
- *Post to all buyers price $p = \mu_B$.*

Proof. The proof is similar to that of Theorem 9, but certain points need some special care. Let κ again be the maximum number of sellers that can be matched to distinct buyers that follow them, but this time under the added restriction of the K-size stock. This corresponds to the maximum matching with no "temporal" cut of size greater than K. We write "temporal" cut to mean any cut in the graph that separates the vertices (buyers and sellers) $1 \dots i$ from vertices $i+1 \dots n$ — that is, precisely the condition that we cannot match more than K sellers from an initial segment to buyers later in the sequence.

In the full version of our paper we show that such a κ-size matching can be computed not only offline, but also online using a FIFO queue of length K, adding sellers to the queue while it is not full and matching buyers greedily: we post prices to sellers, only if we have free space in our stock, i.e. when the matching queue is not full. We underestimate the online profit by considering only selling an item to the buyer that is matched to the seller from which we bought the item. Mimicking the analysis in the proof of Theorem 9 we can see that the expected number of items bought from the κ matched sellers is $\kappa F_S(q) \geq \kappa \frac{1}{2eK}\frac{1}{r}$.

Now we argue that $q \leq \frac{\mu_B}{2}$. Indeed, since $F_S(q) \leq \frac{1}{e}$ we know for sure that $q \leq \mu_S$, and so from Lemma 2 it is $q \leq e\mu_S F(q) \leq e\mu_S \frac{\mu_B}{\mu_S}\frac{1}{2e} = \frac{\mu_B}{2}$. Next, notice that whenever we make a successful sale, the contribution to profit is $p - q \geq \mu_B - \frac{\mu_B}{2} = \frac{1}{2}\mu_B$.

The rest of the proof can be found in the full version of the paper. ◀

▶ Remark. The above upper bound in Theorem 11, although a substantial improvement from the $\Theta(\sqrt{n})$ one for the general case in Theorem 9, cannot be improved further: the logarithmic lower bound is unavoidable, since a careful inspection of the welfare lower bound in the proof of Lemma 6 reveals that the same analysis carries over to the profit.

6 Balanced Sequences

As we saw in Section 5, introducing a restriction in the size of available stock can improve the performance of our online algorithms with respect to profit. However, the bound is still super-constant. Thus, it is perhaps more reasonable to assume some knowledge of the ratio α between buyers and sellers in sequences the intermediary might face. This allows us finer control over the trade-off between high volume of trades and the hunt for greater order statistics.

In this section we analyse the competitive ratio for profit and welfare obtained by online algorithms on α-*balanced* sequences.

▶ **Definition 12.** Let α be a positive integer. A sequence containing m buyers is called α-balanced if it contains αm sellers and the i-th buyer is preceded by at least αi sellers.

For example, the sequence $SBSSBSBB$ is 1-balanced, but $SBBSSB$ is not. Note that since $n = n_S \frac{\alpha+1}{\alpha} = n_B(\alpha+1)$, we only need to know the number of buyers of a sequence. For convenience, we will denote it by m instead of n_B, as it is used quite often. This constraint eliminates the pathological counterexamples of previous sections (such as SB^m) and introduces a much needed "recurrent" flavour to the market: items are constantly traded and in higher quantities, leading to greater profits for both online and offline algorithms.

6.1 Profit

We first work on profit, deriving bounds for a variety of online and offline mechanisms. Naturally, there are two types of offline mechanisms: adaptive and non-adaptive. The *non-adaptive* posted-price mechanism calculates all prices in advance based on the sequence of buyers and sellers, while the *adaptive* posted-price mechanism can alter the prices on the fly, depending on the outcomes of previous trades.

We show that there is a competitive online mechanism for α-balanced sequences. To do this, we compare the optimal adaptive and non-adaptive profit to the profit of a class of hypothetical mechanisms, called *fractional mechanisms*, which are allowed to buy fractional quantities of items: posting the price p would buy exactly $F_S(p)$ items or sell $1 - F_B(p)$ items. The advantage of using fractional mechanisms is that at any point we know the exact quantity of items in the hands of the intermediary instead of the expectation; an immediate consequence of this is that we know in advance whether there is enough quantity to sell, which implies that *the adaptive and non-adaptive versions of the optimal fractional mechanism are identical.*

We can now give an outline of the results in this section: For α-balanced sequences σ with m buyers and αm sellers, we establish the following relations of optimal profits:

$$\text{adaptive}(\sigma) \le \text{fractional}(\sigma) \le \text{fractional}(S^{\alpha m}B^m) \approx \text{non-adaptive}(\sigma), \tag{3}$$

the last of which will be our online algorithm. We begin by the fractional offline mechanism.

▶ **Theorem 13.** *The profit gained by the optimal fractional mechanism for the sequence* $S^{\alpha m}B^m$ *is*

$$\begin{aligned} \max \quad & m\left(p(1 - F_B(p)) - \alpha \cdot qF_S(q)\right) \\ s.t. \quad & 1 - F_B(p) = \alpha F_S(q) \\ & p, q \in [0, \infty). \end{aligned} \tag{4}$$

For other sequences containing αm sellers and m buyers in a different order, we can use the following lemma to establish the middle part of inequality 3.

▶ **Lemma 14.** *For any α-balanced σ with m buyers, fractional$(\sigma) \le$ fractional$(S^{\alpha m}B^m)$*

▶ **Theorem 15.** *For any sequence σ we have adaptive$(\sigma) \le$ fractional(σ).*

The intuition behind the proof of the theorem is that the optimal adaptive profit is bounded from above by the optimal fractional adaptive profit (since fractional mechanisms is a more general class of mechanisms); since in fractional mechanisms optimal adaptive and non-adaptive profits are the same, the theorem follows. For a more rigorous technical treatment, see the full version of our paper.

At this point, we have a clear model of the adversary's power: the fractional mechanism's revenue for sequence $S^{\alpha m}B^m$, setting only two prices p, q for sellers and buyers. Could we do the same online? It seems likely. After all, long sequences of buyers and sellers seem to lead to a similar amount of trading on average by a mechanism setting the same prices.

Based on the previous discussion we propose the following online posted price algorithm:

- Use prices p, q given by the optimal fractional solution for $S^{\alpha m}B^m$(see Theorem 13).

This algorithm works without knowing the length of the sequence chosen by the adversary.

▶ **Lemma 16.** *Let A be the online algorithm defined by the optimal fractional offline prices of* (4)*. Consider two α-balanced sequences σ_1 and σ_2 of equal length. We write $\sigma_1 \succ \sigma_2$ whenever every prefix of σ_1 contains more sellers than the prefix of σ_2 having equal length. Then, $\sigma_1 \succ \sigma_2 \Rightarrow \mathcal{R}(A, \sigma_1) \geq \mathcal{R}(A, \sigma_2)$*

Although not all sequences are comparable (e.g. $SSBBSB$ and $SBSSBB$), the sequence $(S^\alpha B)^m$ is the bottom element among all α-balanced sequences of length $(\alpha+1)m$. This is trivial, as any balanced sequence must have at least $\lceil \frac{i}{(\alpha+1)/(\alpha)} \rceil$ sellers for any prefix of length i and $(S^\alpha B)^m$ is tight for this bound.

To formalize our intuition of making the same number of trades in the long run, we reformulate our algorithm in the more familiar setting of random walks. Instead of considering agents separately, each "timestep" would be one sub-sequence $S^\alpha B$, giving m steps in total. Thus, we are interested in the random variables Z_i, denoting the items in stock at the end of each step, starting with $Z_0 = 0$. Knowing the algorithm buys $\alpha m F_S(q)$ items in expectation, the expected profit can be given by

$$\mathcal{R}((S^\alpha B)^m) = (\alpha m F_S(q) - \mathbb{E}\left[Z_m\right])(p - q) - \mathbb{E}\left[Z_m\right] q, \tag{5}$$

which is the revenue of the expected number of trades minus the cost of the unsold items.

▶ **Lemma 17.** $\mathbb{E}\left[Z_m\right] \leq \sqrt{2m\alpha^2 \log m}\left(1 - \frac{2}{m}\right) + 2$

Proof. The process Z_i is almost a martingale but not quite: clearly $\mathbb{E}\left[Z_i\right] \leq \alpha m$ for all i and we do have $\mathbb{E}\left[Z_{i+1} | Z_i \geq 1\right] = Z_i$ since the expected change in items after that step is $\alpha F_S(q) - (1 - F_B(p)) = 0$ by Theorem 13 . However, $\mathbb{E}\left[Z_{i+1} | Z_i = 0\right] > Z_i$, by the no short selling assumption.

We can define Y_i in the same probability space, where $Y_0 = 0$, and

$$Y_{i+1} = Y_i + \begin{cases} Z_{i+1} & \text{if } Y_i > 0 \\ -Z_{i+1} & \text{if } Y_i < 0 \\ \begin{cases} Z_{i+1} & \text{with probability } \frac{1}{2} \\ -Z_{i+1} & \text{with probability } \frac{1}{2} \end{cases} & \text{if } Y_i = 0 \end{cases}. \tag{6}$$

The crucial observation is that Y_i behaves similar to Z_i but has no barrier at 0. Notice, that $|Y_i| \geq Z_i$ for all i and Y_i is a martingale.

Moreover, we have that $|Y_{i+1} - Y_i| \leq \alpha$ thus by the Azuma-Hoeffding inequality we can bound the expected value $\mathbb{E}\left[Z_m\right]$:

$$\Pr[Z_m \geq x] \leq \Pr[|Y_m| \geq x] = \Pr[|Y_m - Y_0| \geq x] \leq 2e^{\frac{-x^2}{2m\alpha^2}} \Rightarrow \tag{7}$$

$$\mathbb{E}\left[Z_m\right] \leq x\left(1 - 2e^{\frac{-x^2}{2m\alpha^2}}\right) + 2\alpha m e^{\frac{-x^2}{2m\alpha^2}}, \tag{8}$$

where we can set $x = \sqrt{2m\alpha^2 \log m}$ to obtain the simpler form:

$$\mathbb{E}\left[Z_m\right] \leq \sqrt{2m\alpha^2 \log m}\left(1 - \frac{2}{m}\right) + 2\alpha. \tag{9}$$

◀

▶ **Lemma 18.** *Let $r = \max\left\{2, \frac{\mu_S}{\mu_B}\right\}$. The optimal value of Programme* (4) *is at least $m\frac{\mu_B}{2er}$. Furthermore, at any optimal solution the buyer price has to be at most $p \leq 4\ln(4er)\mu_B$.*

▶ **Theorem 19.** *Under our regularity assumptions, the proposed non-adaptive online mechanism is $(1 + o(\alpha^{3/2} r \log r))$-competitive for any balanced sequence, where $r = \max\left\{2, \frac{\mu_S}{\mu_B}\right\}$.*

Proof. Plugging (9) into (5), we get:

$$\begin{aligned}\mathcal{R}((S^\alpha B)^m) &\geq \alpha m F_S(q)(p-q) - \mathbb{E}\left[Z_m\right](p-q) - \mathbb{E}\left[Z_m\right] q \\ &\geq \alpha m F_S(q)(p-q) - \left(\sqrt{2m\alpha^2 \log m}\left(1 - \frac{2}{m}\right) + 2\alpha\right) p \\ &\geq \alpha m F_S(q)(p-q) - O(\alpha\sqrt{m \ln m} p). \end{aligned} \tag{10}$$

Using Lemma 14, Theorem 15 and Theorem 13 we know that for every α-balanced sequence, the profit of our non-adaptive online algorithm is at least $\mathcal{R}((S^\alpha B)^m)$ and the optimal offline is at most that of the fractional on sequence $S^{\alpha m} B^m$, i.e. $\alpha m F_S(q)(p-q)$. Thus, the second term in (10) bounds the additive difference of the online and optimal offline profit, and its ratio with respect to the offline profit is upper bounded by

$$O\left(\frac{\alpha\sqrt{m \ln m} p}{\alpha m F_S(q)(p-q)}\right) = O\left(\frac{\alpha\sqrt{m \ln m}\mu_B \ln(4er)}{m\frac{\mu_B}{2er}}\right) = O\left(\alpha^{3/2}\sqrt{\frac{\ln n}{n}} r \log r\right).$$ ◀

▶ Remark. Among all 1-balanced sequences, the sequence that gives the maximum profit is not $S^m B^m$; intuitively, by moving buyers earlier in the sequence, we obtain more profit by adapting the remaining buying prices to the outcome of these potential trades. For example, the sequence $S^{m/2} B S^{m/2} B^{m-1}$ has better adaptive profit than the sequence $S^m B^m$ for large m. Our work above shows that the difference is asymptotically insignificant, but it remains an intriguing question to determine the balanced sequence with the maximum profit.

6.2 Welfare

Welfare on balanced sequences also improves the competitive ratio of Theorem 4 to a constant. Intuitively, the reason is that the high volume of possible trades dampens the advantage the adversary has in obtaining higher order statistics from buyers.

▶ **Theorem 20.** *The online auction that posts to any seller and buyer the median of their distribution is* 4*-competitive.*

Notice the above theorem holds without any regularity assumption on the agent value distributions.

Acknowledgements. We want to thank Matthias Gerstgrasser for many helpful discussions and his assistance during the initial development of our paper.

References

1 Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. Posting Prices with Unknown Distributions. In *Innovations in Computer Science (ICS)*, jan 2011. URL: `http://conference.itcs.tsinghua.edu.cn/ICS2011/content/paper/35.pdf`.

2 Moshe Babaioff, Shaddin Dughmi, Robert D. Kleinberg, and Aleksandrs Slivkins. Dynamic pricing with limited supply. *ACM Trans. Economics and Comput.*, 3(1):4, 2015. `doi:10.1145/2559152`.

3 Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Online auctions and generalized secretary problems. *SIGecom Exch.*, 7(2):7:1–7:11, June 2008. `doi:10.1145/1399589.1399596`.

4 Ashwinkumar Badanidiyuru, Robert Kleinberg, and Yaron Singer. Learning on a budget: posted price mechanisms for online procurement. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 128–145. ACM, 2012. URL: `http://dl.acm.org/citation.cfm?id=2229026`.

5 Ziv Bar-Yossef, Kirsten Hildrum, and Felix Wu. Incentive-compatible online auctions for digital goods. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'02, pages 964–970, Philadelphia, PA, USA, 2002. Society for Industrial and Applied Mathematics. URL: `http://dl.acm.org/citation.cfm?id=545381.545506`.

6 Richard E. Barlow and Albert W. Marshall. Bounds for Distributions with Monotone Hazard Rate, I. *The Annals of Mathematical Statistics*, 35(3):1234–1257, sep 1964. URL: `http://projecteuclid.org/euclid.aoms/1177703281`, `doi:10.1214/aoms/1177703281`.

7 Avrim Blum and Jason D. Hartline. Near-optimal online auctions. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1156–1163. Society for Industrial and Applied Mathematics, 2005. URL: `http://dl.acm.org/citation.cfm?id=1070597`.

8 Avrim Blum, Tuomas Sandholm, and Martin Zinkevich. Online algorithms for market clearing. *Journal of the ACM (JACM)*, 53(5):845–879, 2006. URL: `http://dl.acm.org/citation.cfm?id=1183913`.

9 Liad Blumrosen and Shahar Dobzinski. (Almost) Efficient Mechanisms for Bilateral Trading. *arXiv preprint arXiv:1604.04876*, 2016. URL: `http://arxiv.org/abs/1604.04876`.

10 Liad Blumrosen and Thomas Holenstein. Posted prices vs. negotiations: An asymptotic analysis. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, EC'08, pages 49–49, New York, NY, USA, 2008. ACM. `doi:10.1145/1386790.1386801`.

11 Liad Blumrosen and Yehonatan Mizrahi. Approximating gains-from-trade in bilateral trading. In *Web and Internet Economics*, pages 400–413. Springer, Berlin, Heidelberg, December 2016. `doi:10.1007/978-3-662-54110-4_28`.

12 Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.

13 Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 311–320. ACM, 2010. URL: `http://dl.acm.org/citation.cfm?id=1806733`.

14 Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1424–1443. Society for Industrial and Applied Mathematics, 2016.

15 X Deng, P Goldberg, B Tang, and J Zhang. Revenue maximization in a bayesian double auction market. *Theoretical Computer Science*, 2014. `doi:10.1016/j.tcs.2014.04.013`.

16 Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 123–135. Society for Industrial and Applied Mathematics, 2015.

17 Matthias Gerstgrasser, Paul W Goldberg, and Elias Koutsoupias. Revenue maximization for market intermediation with correlated priors. In *International Symposium on Algorithmic Game Theory*, pages 273–285. Springer, 2016.

18 Yiannis Giannakopoulos, Elias Koutsoupias, and Philip Lazos. Online market intermediation. *CoRR*, abs/1703.09279, 2017. URL: `http://arxiv.org/abs/1703.09279`.

19 Yiannis Giannakopoulos and Maria Kyropoulou. The VCG Mechanism for Bayesian Scheduling. In *Proceedings of the 11th Conference on Web and Internet Economics*, volume 9470 of *Lecture Notes in Computer Science*, pages 343–356. Springer, 2015. URL: `http://arxiv.org/abs/1509.07455`, `doi:10.1007/978-3-662-48995-6_25`.

20 Mohammad T. Hajiaghayi. Online auctions with re-usable goods. In *Proceedings of the 6th ACM conference on Electronic commerce*, pages 165–174. ACM, 2005. URL: `http://dl.acm.org/citation.cfm?id=1064027`.

21 Mohammad Taghi Hajiaghayi, Robert Kleinberg, and David C. Parkes. Adaptive limited-supply online auctions. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 71–80. ACM, 2004. URL: `http://dl.acm.org/citation.cfm?id=988784`.

22 Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, volume 7, pages 58–65, 2007.

23 Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, FOCS'03, pages 594–, Washington, DC, USA, 2003. IEEE Computer Society. URL: `http://dl.acm.org/citation.cfm?id=946243.946352`.

24 Elias Koutsoupias and George Pierrakos. On the competitive ratio of online sampling auctions. *ACM Transactions on Economics and Computation*, 1(2):10, May 2013. URL: `http://dl.acm.org/citation.cfm?id=2465769.2465775`.

25 R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450, 1992. `doi:10.1016/0022-0531(92)90091-U`.

26 Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983.

27 Rad Niazadeh, Yang Yuan, and Robert Kleinberg. Simple and near-optimal mechanisms for market intermediation. In *International Conference on Web and Internet Economics*, pages 386–399. Springer, 2014.

28 David C. Parkes. Online mechanisms. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 16. Cambridge University Press, New York, NY, USA, 2007.

29 Qiqi Yan. Mechanism design via correlation gap. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'11, pages 710–719. SIAM, 2011. URL: `http://dl.acm.org/citation.cfm?id=2133036.2133092`.