

Competitive Online Optimization under Inventory Constraints

Qiulin Lin, Hanling Yi*
Information Engineering
The Chinese University of Hong Kong

John Pang
Computing and Math. Sciences
California Institute of Technology

Minghua Chen†
Information Engineering
The Chinese University of Hong Kong

Adam Wierman
Computing and Math. Sciences
California Institute of Technology

Michael Honig
Electrical Engineering
Northwestern University

Yuanzhang Xiao
Hawaii Center for Adv. Comm.
University of Hawaii at Manoa

ABSTRACT

This paper studies online optimization under inventory (budget) constraints. While online optimization is a well-studied topic, versions with inventory constraints have proven difficult. We consider a formulation of inventory-constrained optimization that is a generalization of the classic one-way trading problem and has a wide range of applications. We present a new algorithmic framework, CR-Pursuit, and prove that it achieves the minimal competitive ratio among all deterministic algorithms (up to a problem-dependent constant factor) for inventory-constrained online optimization. Our algorithm and its analysis not only simplify and unify the state-of-the-art results for the standard one-way trading problem, but they also establish novel bounds for generalizations including concave revenue functions. For example, for one-way trading with price elasticity, the CR-Pursuit algorithm achieves a competitive ratio that is within a small additive constant (i.e., $1/3$) to the lower bound of $\ln \theta + 1$, where θ is the ratio between the maximum and minimum base prices.

CCS CONCEPTS

• **Theory of computation** → **Online algorithms**; *Design and analysis of algorithms*; • **Applied computing** → *Decision analysis*;

KEYWORDS

Inventory Constraints; Revenue Maximization; Online Algorithms; One-way Trading; Price Elasticity

ACM Reference Format:

Qiulin Lin, Hanling Yi, John Pang, Minghua Chen, Adam Wierman, Michael Honig, and Yuanzhang Xiao. 2019. Competitive Online Optimization under Inventory Constraints. In *Proceedings of ACM Sigmetrics conference*. ACM, New York, NY, USA, Article 1, 15 pages. https://doi.org/10.475/123_4

*The first two authors contribute equally to the work.

†Corresponding author.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ACM Sigmetrics conference, June 2019, Phoenix, AZ USA

© 2019 Copyright held by the owner/author(s).

ACM ISBN 123-4567-24-567/08/06.

https://doi.org/10.475/123_4

1 INTRODUCTION

Online optimization is a foundational topic in a variety of communities, from machine learning to control theory to operations research. There is a large and active community studying online optimization in a wide range of settings, both looking at theoretical analysis and real-world applications. The applications of online optimization are wide ranging, e.g., multi-armed bandits [7, 15, 51], network optimization (with packing constraints) [31, 32], data center capacity management [39, 43, 49], smart grid control [41, 47, 57], and beyond. Further, a diverse set of algorithmic frameworks have been developed for online optimization, from the use of classical potential functions, e.g., [1, 30], to primal-dual techniques, e.g., [31, 53], to approaches based on receding horizon control, e.g., [44, 45]. Additionally, many variations of online optimization have been studied, e.g., online optimization with switching costs [12, 38, 40], online optimization with predictions [18, 38, 41], convex body chasing [5, 10, 28], and more.

In this paper, we focus on an important class of online optimization problems that has proven challenging: *online optimization under inventory (budget) constraints (OOIC)*. In these problems a decision maker has a fixed amount of inventory, e.g., airlines selling flight tickets or battery owners participating in power contingency reserves market, and must make a decision in each of the T rounds with the goal of optimizing per-round revenue functions. The challenge is that the decision maker does not have knowledge of future revenue functions or when the final round will occur, i.e., the value of T . Further, the strict inventory constraint means that an action now has consequences for future rounds. As a result of this entanglement, positive results have only been possible for inventory constrained online optimization in special cases to this point, e.g., the one-way trading problem [23].

More formally, a decision maker in an OOIC participates in T rounds, without knowing T ahead of time. In each round, the decision maker selects an action $v_t \geq 0$, e.g., an amount to sell, after observing a concave revenue function $g_t(\cdot)$. Though the decision maker observes the revenue function each round before choosing an action, it is typically not desirable to choose an action to maximize the revenue in each round due to the limited inventory Δ . Specifically, the actions are constrained by $\sum_{t=1}^T v_t \leq \Delta$, and consequently an action taken at time t constrains future actions. In particular, if the inventory is used too early then better revenue functions may appear later, when inventory is no longer available.

OOIC generalizes many well-known online learning and revenue maximization problems. One of the most prominent is the one-way

trading problem [23], where a trader owns some assets (e.g., dollars) and aims to exchange them into other assets (e.g., yen) as much as possible, depending on the price (e.g., exchange rate). There is a long history of work on one-way trading [19, 21, 23, 29, 42, 56], as we describe in Sec. 7, and OOIC includes both the classic one-way trading problem and variations with concave revenue functions and price elasticity.

Applications. Beyond the one-way trading problem, OOIC also captures a variety of other applications. Three examples that have motivated our interest in OOIC are (i) power contingency reserve markets [2, 52], (ii) network spectrum trading [13, 48], and (iii) online advertisement.

In power contingency reserve markets, the system operator faces a contingency, e.g., shortfall of supply that may lead to cascading blackouts, and communicates this need to either supplement the power system using battery or cut down large scale power supply. Consider the perspective of a battery supply owner that is deciding when to take part in a contingency. A contingency may be solved immediately, or it may instead cause a larger contingency whereby the system operator is willing to pay more at a later time epoch. In preparation to participate in these contingencies, batteries are charged earlier and therefore the marginal cost of participation only manifests as an opportunity cost against future participation in the day. These situations highlight the need for the online properties considered in our work: (i) the unknown ending time T , (ii) future revenue functions are not known, and (iii) a costless, strict inventory constraint.

Similarly, in spectrum trading, the owner of a spectrum band sells bandwidth to make sure that profit or revenue is maximized given the investments that have already been made to procure the particular bandwidth. This means that any cost with regards to sales only appears as opportunity cost against future possible sales. Similarly, a potential buyer who is turned down may seek bandwidth from a different provider, and may never return, or situations may change between time epochs, highlighting the same three properties as before: (i) the unknown ending time T , (ii) future revenue functions are not known, and (iii) a costless and strict inventory constraint.

In online advertisement, an advertiser with a given budget would like to invest into keywords from Internet search engines, e.g., Google AdWords. Potential keywords come in an online fashion and may be unavailable at any time. It has also been shown in [22] that revenue can be modelled as a concave function with respect to the investment. The advertiser needs to decide how to invest its budget for keywords to maximize the overall revenue, once again highlighting the same three properties listed above.

Contributions. In this paper we develop a new algorithmic framework, called CR-Pursuit, and apply it to develop online algorithms for the OOIC problem with the optimal competitive ratio (up to a problem-dependent constant factor). Further, we prove that CR-Pursuit provides the first positive results for a generalization of the classic one-way trading problem with concave revenue functions and price elasticity. In more detail, we summarize our contributions as follows.

First, we introduce a new algorithmic framework, CR-Pursuit, in Sec. 5. The framework is based on the idea of “pursuing” an

optimized competitive ratio at all time. The framework is parameterized by a tight upper bound on the competitive ratio, which is then “pursued” with the actions in each round. We apply the framework to OOIC and generalizations of the one-way trading problem in this paper, but the framework has the potential for broad applicability beyond these settings as well. Along the way, we also derive several useful results on the offline optimal solution in Sec. 4, which may be of independent interest.

Second, in Sec. 6, we apply CR-Pursuit to the OOIC problem to achieve the optimal competitive ratio among all deterministic algorithms (up to a problem-dependent constant factor). To obtain these bounds we use two technical ideas that are of general interest beyond OOIC. First, we prove that it suffices to focus on the single-parametric CR-Pursuit algorithm for achieving the optimal competitive ratio, thus significantly reducing the search space of optimal online algorithms. Second, we identify a “critical” input sequence that highlights an important structural property of the space of input sequences. By applying CR-Pursuit to this critical sequence, we characterize a lower bound on the optimal competitive ratio as $\ln \theta + 1$ where θ is the ratio between the maximum and minimum base prices to be defined in Sec. 3. Subsequently, for any other input, the performance ratio achieved by CR-Pursuit is upper bounded by the product of a problem-dependent factor and the lower bound. This structure not only suggests a principled approach to characterizing the optimal competitive ratio, but also immediately shows that CR-Pursuit achieves the optimal competitive ratio (up to a problem-dependent factor) among all deterministic algorithms.

Third, we apply CR-Pursuit to one-way trading problems in Sec. 7. The novel framework simplifies and unifies the state-of-the-art results of the classic one-way trading problem. In particular, the critical input discussed above is simply the worst case one for classical one-way trading; hence, CR-Pursuit achieves the optimal competitive ratio $\ln \theta + 1$. Further, we show that CR-Pursuit performs well for generalizations of one-way trading where no positive results were previously known. Specifically, for one-way trading with price elasticity and concave revenue functions, CR-Pursuit achieves a competitive ratio that is within a small additive constant (i.e., $1/3$) to the general lower bound of $\ln \theta + 1$.

2 RELATED WORK.

Online optimization is a large and rich research area and excellent surveys can be found in [3, 25]. Well-known problems in the online optimization paradigm include the classic secretary problem [20], the ski rental problem [36], the one-way trading problem [23], and the k -server problem [27]. Our results represent the most general results to date for a situation where actions are subject to a fixed inventory constraint.

The problem considered here is a generalization of the classical one-way trading problem, which has received considerable attention, e.g., [19, 21, 23, 29, 42, 56]. In the one-way trading problem an online decision maker is sequentially presented with exchange rates within a bounded region, and she desires to trade all her assets to another. The amount of assets traded in a single time period is assumed to be small enough to not affect the eventual price. El-Yaniv et al. [23] propose a threshold-based online algorithm with competitive ratio $O(\ln \theta)$. Any remaining items must be sold at

the last epoch as that is revenue maximizing. On the other hand, our analysis allows for leftover inventory (since selling all assets at the last time step may not be revenue maximizing solution for the last time step in the presence of price elasticity or concave revenue functions) and an unknown stopping time, while retaining the competitive ratio.

Variants of the one-way trading problem have been studied in the literature. Chin et al. [19] and Damaschke et al. [21] study the one-way trading problem with unbounded prices and time-varying price bounds, respectively. Zhang et al. [56] study the problem when every two consecutive prices are interrelated. Fujiwara et al. [29] study the problem using average-case competitive analysis under the assumption that the distribution of the maximum exchange rate is known. Kakade et al. [35] incorporate market volume information and study another one-way trading model in the stock market, called the price-volume trading problem. While the classical one-way trading problem mostly deals with linear revenue functions, we note that in our problem we consider general concave revenue functions, which allows us to capture a boarder class of interesting settings, e.g., one-way trading with price elasticity.

Beyond the one-way trading problem, OOIC is also highly related to generalizations of the secretary problem and prophet inequalities, e.g., [9, 24, 50]. Strong positive results have been obtained for these problems; however the analytic setting considered differs dramatically from the current paper. Specifically, we consider a worst case analysis whereas analysis of the secretary problem and prophet inequalities focus on stochastic instances. Under the stochastic setting, so-called “thresholding” algorithms are effective; however such algorithms have unbounded competitive ratios in the worst case setting, even under the simplest assumptions.

Prior to this work, the most general results known for online problems with inventory constraints are for the class of problems termed online optimization with packing constraints, e.g., [6, 8, 11, 16, 17]. This stream of work developed an interesting algorithmic framework based on a primal-dual or multiplicative weights update approaches, which centers around maintaining a dual variable for each constraint, understood as a shadow (or pseudo) price for the constraint given the information thus far. While the inventory constraints we consider are packing constraints, our formulation is fundamentally different than the formulation considered in these papers. In these papers, the constraints come in an online fashion; whereas in our work, the revenue functions arrive in an online fashion.

Another related online optimization problem is the k -search problem, where a player searches for the k highest prices in a sequence that is revealed to her sequentially. When $k \rightarrow \infty$, the k -max search problem becomes the one-way trading problem [42]. Lorenz et. al. [42] propose optimal deterministic and randomized online algorithms for both the k -max search and k -min search problem. This is different from the well-known k -server problem, where an online algorithm must control the movement of k servers in a metric space to minimize the movement (or latency involved) in serving future requests. A popular algorithmic framework for the k -server problem is the *potential function framework*. In contrast to our CR-Pursuit approach, the potential function approach requires a bound between the offline optimal cost and the online cost at each time epoch with respect to the potential.

Finally, it is important to distinguish our work from the literature studying *regret* in online optimization, e.g., [18, 34]. While regret is a natural measure for many online optimization problems, when inventory constraints are present it is no longer appropriate to compare against the best static action, as is done by regret. Static actions are poor choices when optimizing revenue subject to inventory constraints. Instead, competitive ratio is the most appropriate measure. Further, note that there is a fundamental algorithmic trade-off between optimizing regret and competitive ratio, even when inventory constraints are not present. In particular, [4] shows that no algorithm can obtain both sub-linear regret and constant competitive ratio.

3 PROBLEM FORMULATION

We study an online optimization problem where a decision maker sells inventory across an interval of discrete time slots in order to maximize the aggregate revenue. The revenue functions of individual slots are revealed sequentially in an online fashion, and the interval length is unknown to the decision maker. The initial inventory is given in advance as a constraint, and it bounds the decision maker’s aggregate selling quantities across time slots.¹ The key notations used in this paper are summarized in Tab. 1. Throughout this paper, we use $[n]$ to represent the set $\{1, 2, \dots, n\}$ where n is a positive integer.

Table 1: Summary of Notations.

T	The number of time slots
Δ	The initial inventory
$g_t(v)$	The revenue function of time slot t
$\sigma^{[1:t]}$	Input (revenue function) sequence up to time t , i.e., $\{g_1, g_2, \dots, g_t\}$
$p(t)$	Base price at time t , i.e., $g'_t(0)$
m, M	The lower and upper bounds of $p(t)$, $\forall t \in [T]$
θ	The ratio of M/m
λ	The dual variable associated with the inventory constraint in OOIC
v_t	The selling quantity at time t
\bar{v}_t	The selling quantity of CR-Pursuit(π) at time t
v_t^*	The optimal selling quantity at time t under the offline setting
\hat{v}_t	A maximizer of $g_t(v)$ over $[0, \Delta]$
$\Phi_\Delta(\pi)$	The worst case (maximal) inventory over all possible sequences of inputs needed to maintain a competitive ratio $\pi \geq 1$ for CR-Pursuit(π)

More specifically, at time $t \in [T]$, upon observing the revenue function $g_t(\cdot)$, the decision maker has to make an irrevocable decision on an action (quantity) v_t . Upon choosing v_t the decision

¹We emphasize that in contrast to the works on online optimization with packing constraints [6, 8, 11, 16, 17], the uncertainty in our optimization problem is not that the inventory constraint is unknown beforehand, but rather the revenue functions arrives in an online fashion.

maker receives a revenue of $g_t(v_t)$. The overall objective is to maximize the aggregate revenue, while respecting the inventory constraint $\sum_{t \in [T]} v_t \leq \Delta$.² We assume that $g_t(\cdot), \forall t \in [T]$, satisfy the following conditions:

- $g_t(v)$ is concave, increasing, and differentiable over $[0, \Delta]$;
- $g_t(0) = 0$;
- $p(t) \triangleq g'_t(0) > 0$ and $p(t) \in [m, M]$.

The first condition is a smoothness condition on the revenue function and a natural diminishing return assumption. It also limits our discussion in the more interesting setting where at each time, selling more could never decrease revenue. The second condition implies that selling nothing yields no revenue. The third condition limits the marginal revenue at the origin (named base price hereafter) and ensures that it is beneficial to sell, since the base price is positive. Denote the family of all possible revenue functions at time t as \mathcal{G} . We assume m and M are known beforehand to the decision maker and denote $\theta = M/m$.

We formulate the problem of online optimization under inventory constraints (OOIC) as follows:

$$\begin{aligned} \text{OOIC : } \quad & \max \quad \sum_{t=1}^T g_t(v_t) & (1) \\ & \text{s.t.} \quad \sum_{t=1}^T v_t \leq \Delta, & (2) \\ & \text{var.} \quad v_t \geq 0, \forall t \in [T]. & (3) \end{aligned}$$

Without loss of generality, we assume that the inventory constraint in (2) is active at the optimal solution.

We can interpret the inventory constraint (2) in an OOIC in a parallel way to the inventory constraint in the one-way trading problem [23]. In particular, in the one-way trading problem the trader has to decide in each slot the selling quantity v_t to maximize the total revenue at the stopping time T . In fact, when setting the family of functions \mathcal{G} to be the family of revenue functions of the form $g_t(v_t) = p(t)v_t$, we can see OOIC covers the one-way trading problem as a special case. Additionally, when addressing revenue functions of the form $g_t(v_t) = v_t(p(t) - f_t(v_t))$ where f_t is a convex function representing price elasticity, OOIC represents a generalized one-way trading problem with price elasticity.

To study the performance of an online algorithm for OOIC we use the *competitive ratio* as the metric of interest.³ Let \mathcal{A} be a deterministic online algorithm. It is called π -*competitive* if

$$\pi = \max_{\sigma \in \Sigma} \frac{\eta_{OPT}(\sigma)}{\eta_{\mathcal{A}}(\sigma)},$$

²The assumption that the action is chosen *after* observing the function differs from the classical online convex optimization literature [34, 37], but matches the literature on online convex optimization with switching costs [12, 38, 40] and the literature on competitive algorithm design, including those on buy-or-rent decision making problems [36, 43, 57] and metrical task systems [14, 26, 41]. It allows an isolation of the inefficiency resulting from inventory constraints rather than also including the inefficiency resulting from the lack of knowledge of the function.

³Note that many papers in the online optimization literature, e.g., [34], focus on *regret* instead of competitive ratio, but regret is not an appropriate measure when inventory constraints are considered since static actions are no longer appropriate. Our focus on competitive ratio matches that of the literature on secretary problems [9, 50], prophet inequalities [33, 50], online optimization with switching costs [12, 38, 40, 41, 43], etc.

where Σ is the set of all possible inputs ($g_t(\cdot), t \in [T]$) and $\eta_{OPT}(\sigma)$ and $\eta_{\mathcal{A}}(\sigma)$ are the revenues generated by the optimal offline algorithm OPT and the online algorithm \mathcal{A} , respectively. This value π is the competitive ratio (CR) of the algorithm \mathcal{A} .

4 INSIGHTS ON THE OFFLINE SOLUTION

In this section, we derive several results on the optimal offline solution. They are useful in the design and analysis of our algorithmic framework CR-Pursuit in Sec. 5.

Under the offline setting where $g_t(\cdot), \forall t \in [T]$, are known in advance to the decision maker, OOIC is a convex problem and can be solved efficiently. Let v^* be the optimal primal solution and λ^* be the optimal dual variable associated with the inventory constraint in (2). We note that λ^* can be obtained by the algorithm in Alg. 2 in Appendix A.1, based on a binary search idea. The following proposition gives a set of optimality conditions for the optimal primal solutions v^* and the optimal dual variable λ^* .

Proposition 1. *Under our setting that the inventory constraint is active at the optimal solution, the optimal primal and dual solutions v^* and λ^* satisfy (i) $\lambda^* \geq 0$ and $\sum_{t=1}^T v_t^* = \Delta$ and (ii) for each $t \in [T]$,*

$$\begin{cases} v_t^* = 0, & \text{if } g'_t(0) < \lambda^*; \\ v_t^* \in V_t(\lambda^*) \triangleq \{v_t | g'_t(v_t) = \lambda^*, v_t \in [0, \Delta]\}, & \text{otherwise.} \end{cases} \quad (4)$$

Recall that at time t , the marginal revenue evaluated at v_t is $g'_t(v_t)$, which is no larger than the base price $p(t) = g'_t(0)$ due to the concavity of $g_t(\cdot)$. The optimal dual variable λ^* can be interpreted as the marginal cost (shadow price) of the inventory. Then Proposition 1 says that, at the optimal solution, the marginal revenue must equal the marginal cost in the slots with positive selling quantities. Moreover, it is optimal to sell only in the slots in which the base price is higher than the optimal marginal cost, i.e., $p(t) > \lambda^*$. These observations are similar to those in the Cournot competition literature, e.g., [46].

Next, we reveal two interesting observations on the offline optimal aggregate revenue. Recall the input (revenue function) sequence until time t is

$$\sigma^{[1:t]} = \sigma^{[1:t-1]} \cup \{g_t(\cdot)\} = \sigma^{[1:t-2]} \cup \{g_{t-1}(\cdot)\} \cup \{g_t(\cdot)\} = \dots$$

Recall that $\eta_{OPT}(\sigma^{[1:t]})$ is the offline optimal aggregate revenue given the input $\sigma^{[1:t]}$. The following lemma bounds the increment of the optimal aggregate revenue as t increases.

Lemma 2. *Let λ_{t-1} and λ_t be the optimal dual variables associated with the inventory constraint given the inputs $\sigma^{[1:t-1]}$ and $\sigma^{[1:t]} = \sigma^{[1:t-1]} \cup \{g_t(\cdot)\}$, respectively. Let \tilde{v}_t be the optimal offline solution in the (last) time slot t given the input $\sigma^{[1:t]}$. The following inequalities hold:*

$$\eta_{OPT}(\sigma^{[1:t]}) - \eta_{OPT}(\sigma^{[1:t-1]}) \geq g_t(\tilde{v}_t) - \lambda_t \tilde{v}_t, \quad (5)$$

and

$$\eta_{OPT}(\sigma^{[1:t]}) - \eta_{OPT}(\sigma^{[1:t-1]}) \leq g_t(\hat{v}_t) - \lambda_{t-1} \hat{v}_t \leq g_t(\hat{v}_t), \quad (6)$$

where \hat{v}_t is the maximizer of $g_t(\cdot)$ over $[0, \Delta]$.

Note that λ_{t-1} and λ_t are the marginal costs of inventory at the optimal solutions to OOIC given the inputs $\sigma^{[1:t-1]}$ and $\sigma^{[1:t]} = \sigma^{[1:t-1]} \cup \{g_t(\cdot)\}$, respectively. The terms $\lambda_t \tilde{v}_t$ and $\lambda_{t-1} \tilde{v}_t$ represent the upper and lower bounds of the cost of committing \tilde{v}_t to obtain the new additional revenue $g_t(\tilde{v}_t)$ in time slot t . Thus the difference between them represents a bound on the “profit” obtained in slot t . Intuitively, Lemma 2 says that one can bound the optimal offline revenue increment by these profit bounds, as shown in (5) and (6).

The proof of Lemma 2 is included in the Appendix A.3; we give the proof idea here. Compared with the optimal solution under $\sigma^{[1:t-1]}$, the optimal solution under $\sigma^{[1:t]}$ is smaller at τ , $\forall \tau \leq t-1$, in order to commit \tilde{v}_t to $g_t(\cdot)$, which cause a decrement in revenue. Furthermore, the per-unit revenue lost is upper bounded by λ_t and lower bounded by λ_{t-1} . Combining the two understandings gives the bounds of the increment on optimal revenue at each time.

The upper bound in (6) also highlights an intuitive result that the increment of the optimal aggregate revenue from $t-1$ to t is at most $g_t(\tilde{v}_t)$, i.e., the maximum revenue one can obtain in slot t .

Our last result in this section, as stated in the lemma below, reveals another subtle yet important property of the increment of the optimal aggregate revenue.

Lemma 3. *Let $\tilde{\sigma}$ be an input sequence. $\bar{\sigma}$ is another input sequence constructed by interchanging g_τ and $g_{\tau+1}$ in $\tilde{\sigma}$, for any selected $\tau \in [T]$. We have*

$$\eta_{OPT}(\tilde{\sigma}^{[1:\tau]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau-1]}) \geq \eta_{OPT}(\bar{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\bar{\sigma}^{[1:\tau]}). \quad (7)$$

The left- (resp. right-) hand-side of (7) can be regarded as the increment g_τ contributes to the offline optimal under $\tilde{\sigma}$ (resp. $\bar{\sigma}$). Inequality (7) means that moving $g_{\tau+1}$ ahead of g_τ (as under $\bar{\sigma}$) will not increase the contribution of g_τ to the offline optimal. Lemma 3 basically states that regardless of the input sequence thus far, the impact or improvement in the offline optimal that g_τ brings at the time it appears in the input sequence has a “diminishing effect” in time. The proof of Lemma 3 is essentially based on the bounds on the increment of offline optimal at each time in Lemma 2. We leave the proof in Appendix A.4.

5 CR-PURSUIT ALGORITHMIC FRAMEWORK

In this section, we present CR-Pursuit, a new algorithmic framework for solving OOIC under the online setting, where the interval length T is not known beforehand and the revenue functions $g_t(\cdot)$, $t \in [T]$, are revealed in a slot-by-slot fashion.

CR-Pursuit is parameterized by a competitive ratio π , and chooses actions with the goal of “pursuing” this competitive ratio, i.e., maintaining the competitive ratio against the offline optimal of the previously observed revenue functions at all time. We derive bounds on the optimal competitive ratios and use them to operate CR-Pursuit accordingly.

In the following, we first present the CR-Pursuit framework and show that one can optimize the only parameter of CR-Pursuit to achieve the best possible competitive ratio, thus significantly reducing the search space of optimal online algorithms. Then, we identify a “critical” input sequence that highlights an important structural property of the space of input sequences. By applying CR-Pursuit to this critical sequence, we characterize a lower bound

Algorithm 1 CR-Pursuit(π) Online algorithm

- 1: **Input:** $\pi > 1, \Delta$
 - 2: **Output:** $\bar{v}_t, t \in [T]$
 - 3: **while** t is not the last slot **do**
 - 4: Obtain $\eta_{OPT}(\sigma^{[1:t]})$ by solving the convex problem OOIC given the input until t , i.e., $\sigma^{[1:t]}$
 - 5: Obtain a $\bar{v}_t \in [0, \Delta]$ that satisfies (8)
 - 6: **end while**
-

on the optimal competitive ratio as $\ln \theta + 1$, where we recall that $\theta = M/m$ is the ratio between the maximum and minimum base prices. Then, for any other input, the performance ratio achieved by CR-Pursuit (with the same parameter) is upper bounded by the product of a problem-dependent factor and the lower bound. This structure not only suggests a principled approach for characterizing the optimal competitive ratio, but also immediately shows that CR-Pursuit (with a parameter being the product of the lower bound and the problem-dependent factor) achieves the optimal competitive ratio (up to a problem-dependent factor) for solving OOIC among all deterministic algorithms.

CR-Pursuit. Recall that $\sigma^{[1:t]} = \{g_1, g_2, \dots, g_t\}$ is the input up to time t and $\eta_{OPT}(\sigma^{[1:t]})$ is the corresponding optimal offline revenue. The class of online algorithms that make up the CR-Pursuit framework, denoted as CR-Pursuit(π) and presented in Alg. 1, can be described as follows: Given any $\pi \geq 1$, at the current time t , CR-Pursuit(π) outputs a $\bar{v}_t \in [0, \Delta]$ that satisfies

$$g_t(\bar{v}_t) = \frac{1}{\pi} \left[\eta_{OPT}(\sigma^{[1:t]}) - \eta_{OPT}(\sigma^{[1:t-1]}) \right]. \quad (8)$$

We remark that such \bar{v}_t always exists, because (i) $g_t(\cdot)$ is a continuous and increasing function and (ii) the right-hand-side of (8) is in $[g_t(0), g_t(\Delta)]$ according to Lemma 2.

Essentially, CR-Pursuit(π) aims at keeping the *offline-to-online revenue ratio* to be $\pi > 1$ at all time, i.e.,

$$\sum_{\tau=1}^t g_\tau(\bar{v}_\tau) = \frac{1}{\pi} \eta_{OPT}(\sigma^{[1:t]}), \quad \forall t \in [T]. \quad (9)$$

While CR-Pursuit(π) can be defined for any π , the solution obtained by CR-Pursuit(π) may violate the inventory constraint in OOIC and be infeasible. This motivates the following definition.

Definition 4. *CR-Pursuit(π) is feasible if $\Phi_\Delta(\pi) \leq \Delta$, where*

$$\Phi_\Delta(\pi) \triangleq \max_{\sigma \in \Sigma} \sum_{t=1}^T \bar{v}_t(\sigma), \quad (10)$$

and $\bar{v}_t(\sigma)$ is the output of CR-Pursuit(π) at time t under the input σ .

If CR-Pursuit(π) is feasible, i.e., it can maintain the offline-to-online revenue ratio to be π under all possible input sequences without violating the inventory constraint, then by definition it is π -competitive. We present a useful observation on $\Phi_\Delta(\pi)$.

Lemma 5. *$\Phi_\Delta(\pi)$ is strictly decreasing in π over $[1, \infty)$.*

Lemma 5 follows naturally since attempting to preserve a smaller competitive ratio requires selling a larger inventory to match the discounted revenue obtained by the offline optimal algorithm. It

also implies that if CR-Pursuit(π_1) is feasible for some π_1 , then any online algorithm CR-Pursuit(π) with $\pi \geq \pi_1$ is also feasible. Thus an upper bound on the optimal competitive ratio in this case gives a feasible competitive online algorithm.

The Optimal Competitive Ratio. We now present a key result, which says that it suffices to focus on CR-Pursuit for achieving the optimal competitive ratio.

Theorem 6. *Let π^* be the unique solution to the characteristic equation $\Phi_\Delta(\pi) = \Delta$. Then CR-Pursuit(π^*) is feasible and π^* is the optimal competitive ratio of deterministic online algorithms.*

Before we proceed to prove Theorem 6, we first present the following lemma characterizing a class of worst case inputs for CR-Pursuit. The results will be used in the proof of Theorem 6.

Lemma 7. *For any CR-Pursuit(π), there exists an input sequence σ such that (i) CR-Pursuit(π) sells exactly the $\Phi_\Delta(\pi)$ amount of inventory and (ii) $g'_t(\bar{v}_t(\sigma))$ is non-decreasing in t .*

Lemma 7 states that to compute $\Phi_\Delta(\pi)$, it is sufficient to focus on the input sequences that will lead to a non-increasing sequence of marginal revenue $g'_t(\bar{v}_t)$ at the solution obtained by CR-Pursuit(π). Intuitively, these sequences will cause the CR-Pursuit(π) algorithm to sell large quantities at lower prices in the early slots, without knowing that the marginal revenues at later slots are higher, which is exploited by the offline optimal solution. As a result, CR-Pursuit(π) will need to sell the “worst” amount of inventory to keep the revenue ratio π .

The proof of Lemma 7 is provided in Appendix A.5, based on the subtle yet important property of the offline optimal aggregate revenue in Lemma 3. The idea of the proof, roughly speaking, is that if the worst case input sequence is not as stated, then we can swap revenue functions within the sequence to construct a new worst case one that satisfies the conditions. Putting the preceding lemmas together, we are now ready to prove Theorem 6.

PROOF OF THEOREM 6. The feasibility of CR-Pursuit(π^*) is because of the definition of π^* . What remains to be proved is that π^* is the optimal competitive ratio.

Consider an arbitrary deterministic online algorithm different from CR-Pursuit(π^*), denoted as \mathcal{A} . We will show that \mathcal{A} cannot achieve an offline-to-online revenue ratio smaller than π^* over an input sequence that we construct.

Let $\bar{\sigma}^{[1:T]} = \{\bar{g}_1, \bar{g}_2, \dots, \bar{g}_T\}$ be a worst case input sequence of CR-Pursuit(π^*) that satisfies the conditions in Lemma 7. Let \bar{v}_t and $v_t^{\mathcal{A}}$ be the corresponding solutions of CR-Pursuit(π^*) and \mathcal{A} at time t , respectively. We have

- $\sum_{t=1}^T \bar{v}_t = \Phi_\Delta(\pi^*) = \Delta$;
- $\bar{g}'_t(\bar{v}_t)$ is non-decreasing in t .

We now construct an input sequence over which \mathcal{A} cannot achieve an offline-to-online revenue ratio smaller than π^* , by feeding $\bar{g}_1, \bar{g}_2, \dots, \bar{g}_T$ to \mathcal{A} and stop at any time that we need.

We first present \bar{g}_1 to \mathcal{A} in the first slot. If $v_1^{\mathcal{A}} \leq \bar{v}_1$, we stop and set $T = 1$ in this constructed sequence. In this case, we have

$$\bar{g}_1(v_1^{\mathcal{A}}) \leq \bar{g}_1(\bar{v}_1) = \frac{1}{\pi^*} \eta_{OPT}(\bar{\sigma}^{[1:1]}),$$

thus the competitive ratio of \mathcal{A} is at least π^* . Otherwise we have $v_1^{\mathcal{A}} > \bar{v}_1$ and we continue to present \bar{g}_2 to \mathcal{A} in the second slot.

In general, if at time t the total selling quantity of \mathcal{A} so far is no larger than that of CR-Pursuit(π^*), i.e., $\sum_{\tau=1}^t \bar{v}_\tau$, we end the trading period. Otherwise, we continue to the $t + 1$ slot and present \mathcal{A} with the revenue function $\bar{g}_{t+1}(\cdot)$.

Let τ be the earliest slot such that at the end of time τ , the total selling quantity of \mathcal{A} is less than that of CR-Pursuit(π^*). Such τ exists; otherwise, we will have $\sum_{t=1}^T v_t^{\mathcal{A}} > \sum_{t=1}^T \bar{v}_t = \Delta$, which implies that \mathcal{A} is not feasible. Given such $\tau \in [T]$, we have

$$\sum_{\xi=1}^t v_\xi^{\mathcal{A}} > \sum_{\xi=1}^t \bar{v}_\xi, \forall t \in [\tau - 1], \quad (11)$$

$$\text{and } \sum_{\xi=1}^{\tau} v_\xi^{\mathcal{A}} \leq \sum_{\xi=1}^{\tau} \bar{v}_\xi. \quad (12)$$

We now show that, for the input sequence $\bar{\sigma}^{[1:\tau]}$, the aggregate revenue of \mathcal{A} is no larger than that of CR-Pursuit(π^*), i.e.,

$$\sum_{\xi=1}^{\tau} \bar{g}_\xi(v_\xi^{\mathcal{A}}) - \sum_{\xi=1}^{\tau} \bar{g}_\xi(\bar{v}_\xi) \leq 0, \quad (13)$$

which then implies that the online algorithm \mathcal{A} is at best π^* -competitive. By the concavity of $\bar{g}_t(\cdot)$, we have

$$\begin{aligned} & \sum_{\xi=1}^{\tau} \left[\bar{g}_\xi(v_\xi^{\mathcal{A}}) - \bar{g}_\xi(\bar{v}_\xi) \right] \\ & \leq \sum_{\xi=1}^{\tau} \bar{g}'_\xi(\bar{v}_\xi) (v_\xi^{\mathcal{A}} - \bar{v}_\xi) \\ & = \bar{g}'_\tau(\bar{v}_\tau) \left(\sum_{\xi=1}^{\tau} v_\xi^{\mathcal{A}} - \sum_{\xi=1}^{\tau} \bar{v}_\xi \right) \\ & \quad - \sum_{t=1}^{\tau-1} [\bar{g}'_{t+1}(\bar{v}_{t+1}) - \bar{g}'_t(\bar{v}_t)] \left(\sum_{\xi=1}^t v_\xi^{\mathcal{A}} - \sum_{\xi=1}^t \bar{v}_\xi \right). \end{aligned}$$

By (12) and that $\bar{g}'_\tau(\bar{v}_\tau) \geq 0$ as $\bar{g}_\tau(\cdot)$ is an increasing function, the first term in the last line of derivation is non-positive. By (11) and that $\bar{g}'_t(\bar{v}_t)$ is non-decreasing in t , each term in the summation in the last line of derivation is non-negative. As such, the right-hand-side is non-positive and the inequality in (13) holds. \square

6 COMPETITIVE ANALYSIS OF CR-PURSUIT

The results in the previous section highlight a principled approach to construct an optimal online algorithm. Specifically, the first step is to mathematically characterize $\Phi_\Delta(\pi)$. Then we solve the characteristic equation $\Phi_\Delta(\pi) = \Delta$ to obtain the optimal competitive ratio π^* , and CR-Pursuit(π^*) is an optimal online algorithm for solving OOIC. For special cases such as the one-way trading problem [23] where $g_t(v) = p(t) \cdot v$, we can obtain the closed-form expression of $\Phi_\Delta(\pi)$ and compute the optimal competitive ratio (as demonstrated in Sec. 7.1). However, it is difficult to obtain a closed-form expression for general concave revenue functions. Instead, we characterize an upper bound on $\Phi_\Delta(\pi)$, based on which we can give an upper bound on the optimal competitive ratio π^* and consequently a feasible online algorithm.

Before moving to the upper bound though, it is helpful to understand a lower bound on the optimal competitive ratio. For this, we

can simply refer to the literature on one-way trading. In particular, it has been shown that the optimal competitive ratio of the classic one-way trading problem is $\ln \theta + 1$ [23, 54]. Since OOIC covers one-way trading as a special case, the optimal competitive ratio for any online algorithm solving OOIC is lower bounded by $\ln \theta + 1$. Interestingly, it is possible to interpret this bound in the context of the CR-Pursuit framework. In particular, in Sec. 7.1. we identify the worst case input in one-way trading (defined in Sec. 7.1) as a “critical” input sequence, reflecting an interesting structure on the space of input sequences. By applying CR-Pursuit to this sequence, we characterize a lower bound on the optimal competitive ratio as $\ln \theta + 1$.

It turns out that for any other inputs, the performance ratio achieved by CR-Pursuit is upper bounded by the product of a problem-dependent factor and the lower bound $\ln \theta + 1$. This insight leads to the following results.

Theorem 8. *Recall that \mathcal{G} is the set of all possible $g(\cdot)$ and $\hat{v} \in [0, \Delta]$ is the maximizer of $g(\cdot)$. Let $c = \sup_{g \in \mathcal{G}} \frac{g'(0)}{g(\hat{v})/\hat{v}}$, then the optimal competitive ratio π^* satisfies*

$$\ln \theta + 1 \leq \pi^* \leq c(\ln \theta + 1).$$

Theorem 8 characterizes an upper bound on the optimal competitive ratio in the case for general revenue functions g_t , and also implies that CR-Pursuit($c(\ln \theta + 1)$) is feasible and its competitive ratio is $c(\ln \theta + 1)$. Note that c is a constant that depends on the gradient properties (in particular the base price) and the maximizers of the revenue functions⁴. For many interesting problems, this c is bounded and small. For example, for the one-way trading problem where the revenue functions are linear, i.e., $g_t(v) = p(t)v, \forall t \in [T]$, we have $c = 1$. As another example, for the one-way trading with linear price elasticity where the revenue functions are quadratic, i.e., $g_t(v) = (p(t) - \alpha_t v)v, \forall t \in [T]$, we have $c = 2$.

To prove this theorem, we use a sequence of lemmas elaborated as follows. We begin with Lemma 9, which gives an upper bound on the total selling quantity by CR-Pursuit(π) in each time slot to maintain the offline-to-online revenue ratio. Recall that the output of the algorithm CR-Pursuit(π) at slot t , $g_t(\bar{v}_t)$ is given in (8), and $p(t) = g'_t(0)$ is the base price at slot t .

Lemma 9. *For any input sequence σ , we have*

$$\bar{v}_t \leq c \frac{g_t(\bar{v}_t)}{p(t)}, \forall t \in [T].$$

The proof of Lemma 9 is included in Appendix A.6, by leveraging the definition of c and that $g_t(v)$ is an increasing concave function.

Next, we present an interesting result that bounds the contribution to the online revenue in all the slots whose base prices is no higher than any specific threshold.

⁴While c is a constant when the family of revenue functions are fixed, it is indeed true that c could presumably be driven to be infinitely large, e.g., with revenue functions that are concave and increasing. This parameter c can be seen in an economical sense as a comparison between the base price and the average price at the maximizer of the function. Since the former is already bounded in $[m, M]$, we look at the case when the latter is small. These situations are hard to derive any interesting online optimization as the functions require too much commitment even in bad time epochs, and have low average prices. This results in low committed average prices while the offline optimal may eventually not have to participate in these time epochs.

Lemma 10. *For any input sequence $\sigma \in \Sigma$, for any threshold price $p \in [m, M]$, we have*

$$\sum_{\{t: p(t) \leq p\}} g_t(\bar{v}_t) \leq \frac{1}{\pi} p \cdot \Delta.$$

Lemma 10 is intuitive in that the left-hand-side is the online revenue obtained by CR-Pursuit(π) in the slots whose base prices is not higher than p . The right-hand-side is simply the maximum revenue achievable by CR-Pursuit(π) in these slots according to its design. In the proof, we first observe that if $p(t) < p, \forall t \in [T]$, the result is immediate. As for general cases, based on Lemma 3, we can construct new input sequences by moving forward the slots with $p(t) \leq p$ in σ , while increasing the online revenue in the slots that we are interested in. At last, we obtain an input sequence with larger online revenue in these slots, which are now all in the beginning of the input sequence. The total online revenue in them is bounded by $p \cdot \Delta/\pi$. Lemma 10 allows to prove a key step used in the proof of Theorem 8 below.

Lemma 11. *For any input sequence σ , we have*

$$\sum_{t=1}^T \frac{g_t(\bar{v}_t)}{p(t)} \leq \frac{\Delta}{\pi} (\ln \theta + 1). \quad (14)$$

The idea to prove Lemma 11 is to construct an optimization problem, whose optimal objective value bounds the left-hand-side in (14), subject to the constraint from Lemma 10. Then we show the optimal objective value can be further upper bounded by the right-hand-side in (14).

We are now ready to prove Theorem 8.

PROOF OF THEOREM 8. It is clear that CR-Pursuit is at best $(\ln \theta + 1)$ -competitive, as it covers the one-way trading problem as a special case, which has an optimal competitive ratio of $\ln \theta + 1$.

To establish the upper bound, by Lemmas 9 and 11, we observe

$$\begin{aligned} \Phi_{\Delta}(\pi) &= \max_{\sigma \in \Sigma} \sum_{t=1}^T \bar{v}_t \leq \sum_{t=1}^T c \frac{g_t(\bar{v}_t)}{p(t)} \\ &\leq c \frac{\Delta}{\pi} (\ln \theta + 1). \end{aligned}$$

By solving $c \frac{\Delta}{\pi} (\ln \theta + 1) = \Delta$, we get that $\bar{\pi} = c(\ln \theta + 1)$ and $\Phi_{\Delta}(\bar{\pi}) \leq \Delta$. Then according to the definitions, CR-Pursuit($\bar{\pi}$) is feasible and is $\bar{\pi}$ -competitive. Hence, $\bar{\pi}$ is an upper bound for the optimal competitive ratio π^* . \square

Theorem 8 implies that CR-Pursuit achieves the optimal competitive ratio (up to a problem-dependent factor c) among all deterministic online algorithms.

7 APPLICATION TO ONE-WAY TRADING

In this section, we apply CR-Pursuit to the classic one-way trading problem [23] and its generalizations, illustrating that the framework can both match state-of-the-art results for the classic setting and provide new results for generalizations that have previously resisted analysis. In particular, using the CR-Pursuit framework, we obtain an online algorithm matching the optimal competitive ratio $(\ln \theta + 1)$ for the classic one-way trading problem in Proposition 13 and a near-optimal $(\ln \theta + 4/3)$ result for the case with linear price elasticity in

Theorem 16. Furthermore, the algorithmic framework also extends to any convex price elasticity, and yield online algorithms with order-optimal competitive ratio in these cases.

This section also provides an illustration of how the framework can be applied to specific problem domains to obtain tighter competitive ratio upper bounds than the generic ones under general settings. In particular, for one-way trading with linear price elasticity, the upper bound derived from Sec. 6 is $2(\ln \theta + 1)$ while the bound obtained in this section is $\ln \theta + 4/3$.

In Sec. 7.1, we obtain a close-form expression of $\Phi_\Delta(\pi)$ and compute the optimal π^* in this special case. In Sec. 7.2, we show the ease of generalizing the one-way trading problem, to cases where price formation include price elasticity, an aspect that has been left out and desired in the one-way trading community.

7.1 Classic One-way Trading

In the classic one-way trading problem, a trader owns some assets (e.g., one dollar) at the beginning and aims to exchange it into another assets (e.g., yen) as much as possible, depending on the price (e.g., exchange rate). Thus, the one-way trading problem is a special case of the OOIC problem with $g_t(v_t) = p(t)v_t$ for all $t \in [T]$ and the input at time t can be simplified as $p(t)$.

As a direct application, one can obtain from Sec. 6 that the upper bound for the one-way trading problem is $\ln \theta + 1$, which matches the lower bound. Thus, we immediately know that the optimal competitive ratio for one-way trading is $\ln \theta + 1$ and CR-Pursuit($\ln \theta + 1$) is an optimal deterministic online algorithm. In this section, with the aim of demonstrating the possibility of mathematically characterizing $\Phi_\Delta(\pi)$ in specific problems, we first derive a closed-form expression of $\Phi_\Delta(\pi)$, then we obtain the optimal competitive ratio π^* by solving the characteristic equation $\Phi_\Delta(\pi) = \Delta$.

In the classic one-way trading problem, given any input up to time t , denoted as $\sigma^{[1:t]} \triangleq \{p(1), p(2), \dots, p(t)\}$, the optimal offline revenue can be expressed as $\eta_{OPT}(\sigma^{[1:t]}) = \Delta \cdot \max \sigma^{[1:t]}$. Given any $\pi \geq 1$, we focus on CR-Pursuit(π) defined in Sec. 5. At time t , CR-Pursuit(π) sells the amount $\bar{v}_t \in [0, \Delta]$ that satisfies:

$$\bar{v}_t = \frac{1}{\pi \cdot p(t)} \left[\eta_{OPT}(\sigma^{[1:t]}) - \eta_{OPT}(\sigma^{[1:t-1]}) \right]. \quad (15)$$

As discussed, CR-Pursuit(π) aims at keeping the *offline-to-online revenue ratio* to be $\pi > 1$ at all time.

From Sec. 5, we know that if $\Phi_\Delta(\pi) \leq \Delta$, then CR-Pursuit(π) is feasible and it is π -competitive. In the following, our goal is to derive a close-form expression of $\Phi_\Delta(\pi)$.

Observe that at slot t , the selling decision of CR-Pursuit(π^*) can be simplified as

$$\bar{v}_t = \frac{\Delta}{\pi^* \cdot p(t)} \left(\max \sigma^{[1:t]} - \max \sigma^{[1:t-1]} \right). \quad (16)$$

This suggests that CR-Pursuit(π^*) will sell only when the current price is higher than the best price so far. With this observation, we have the following lemma.

Lemma 12. *For CR-Pursuit(π) with $\pi \geq 1$, given any input $\sigma^{[1:T]}$, to compute $\Phi_\Delta(\pi)$, it is sufficient to consider increasing-price sequences.*

Lemma 12 is a corollary of Lemma 7 in that while the marginal prices are determined by the participation of the algorithm in the

latter, it is constant here in the classic one-way trading problem. Lemma 12 can be proved by observing that the revenue of both the offline and online algorithms remain unchanged if the current price is not the highest price so far, and in that case removing this price from the input sequence will not affect the behaviors of both the offline and online algorithms. From Lemma 12, we know that it is sufficient to consider the following increasing price sequence with length $n \leq T$:

$$m \leq p_1 < p_2 < \dots < p_n \leq M. \quad (17)$$

Under the given price sequence, the optimal offline revenue at time $t \in [n]$ can be simplified as

$$\eta_{OPT}(\sigma^{[1:t]}) = p_t \Delta.$$

According to (16), the output of CR-Pursuit(π) is given by

$$\bar{v}_t = \frac{\Delta}{\pi} \frac{p_t - p_{t-1}}{p_t}, \forall t \in [n],$$

where $p_0 = 0$. Then we have

$$\begin{aligned} \Phi_\Delta(\pi) &= \max_{p_1, p_2, \dots, p_n} \sum_{t=1}^n \bar{v}_t \\ &= \max_{p_1, p_2, \dots, p_n} \frac{\Delta}{\pi} \left(1 + \frac{p_2 - p_1}{p_2} + \dots + \frac{p_n - p_{n-1}}{p_n} \right) \\ &\stackrel{(a)}{=} \frac{\Delta}{\pi} \left(1 + \int_m^M \frac{1}{x} dx \right) = \frac{\Delta}{\pi} (1 + \ln \theta), \end{aligned}$$

where (a) holds when the input sequence in (17) satisfies $n \rightarrow \infty$ and $p_i \rightarrow p_{i+1}$, $\forall i \in [n-1]$. Indeed, this is the worst case input sequence for one-way trading problem, also known as the ‘‘critical’’ input sequence. Setting $\Phi_\Delta(\pi) = \Delta$ yields the solution that $\pi^* = \ln \theta + 1$. Consequently, we have the following result.

Proposition 13. *With $\pi^* = \ln \theta + 1$, CR-Pursuit(π^*) is feasible and an optimal online algorithm for the one-way trading problem.*

The proof follows the same idea and approach as that of Theorem 6. We leave it as an exercise for readers.

7.2 One-way Trading with Price Elasticity

In this subsection, we consider the one-way trading problem in a generalized setting with an additional flexibility on the price model playing the role of *price elasticity*. We assume that price is affected by the total quantity sold at each slot, implying that the decision of how much to sell affects the trading price, usually known in the economics literature as *price elasticity*.

Specifically, we assume that at each slot $t \in [T]$, the price elasticity, defined as $\triangleq f_t(v)$, is a convex non-negative function of the selling quantity with $f(0) = 0$. Under this setting, the revenue function at time t becomes $g_t(v) = (p(t) - f_t(v))v$. This setting can be considered as a special case of OOIC and the input at time t can be simplified as $(p(t), f_t(v))$. Here we have $g'_t(0) = p(t) \in [m, M]$ and $f_t(v) \in [0, +\infty)$, $\forall v \in [0, \Delta]$, $f_t(0) = 0$. Namely, the set of all possible revenue functions can be expressed as

$$\mathcal{G} = \{g_t(v) | g_t(v) = (p(t) - f_t(v))v, p(t) \in [m, M], f_t(v) \in [0, +\infty), \forall v \in [0, \Delta], f_t(0) = 0\}.$$

Note that when $f_t(v) = 0, \forall t \in [T]$, the problem reduces to one-way trading problem considered in Sec. 7.1. Thus we note that any deterministic online algorithm in one-way trading with price elasticity has a competitive ratio of at least $\ln \theta + 1$. When $f_t(v) = \alpha_t v, \alpha_t \geq 0, \forall t \in [T]$, the problem becomes a one-way trading problem with linear price elasticity, which is a common setting in economic literature, e.g., in Cournot competition [46].

Consider the online algorithm CR-Pursuit(π) defined in Sec. 5. When there is price elasticity in the setting, it is difficult to obtain the closed-form expression of $\Phi_\Delta(\pi)$. We follow the analysis in Sec. 6 to obtain an upper bound on $\Phi_\Delta(\pi)$.

In particular, restating Lemma 9 under the parametric assumptions of $g_t(v_t)$ in the one-way trading problem with price elasticity, we can upper bound the selling quantity of CR-Pursuit(π) at each slot with a better characterization of c , reflected in the following lemma.

Lemma 14. *For any input sequence σ , we have*

$$\bar{v}_t \leq c \frac{g_t(\bar{v}_t)}{p(t)}, \forall t \in [T],$$

where

$$c = 2 \left(1 + \sqrt{1 - 1/\pi} \right)^{-1}. \quad (18)$$

We note that value of c given in (18) is smaller than that derived in Lemma 9. The idea of the proof in Appendix A.10 is similar to that of Lemma 9, but we further utilize the special structure of $g_t(\cdot)$ here (i.e., the convexity of $f_t(\cdot)$). The tighter characterization of c allows us to develop an online algorithm with better competitive ratio as compared to the one obtained as a result of Sec. 6.

Lemma 15. *For CR-Pursuit(π) with $\pi \geq 1$, we have*

$$\Phi_\Delta(\pi) \leq \bar{\Phi}(\pi),$$

where $\bar{\Phi}(\pi) \triangleq 2\Delta \left[\pi \left(1 + \sqrt{1 - 1/\pi} \right) \right]^{-1} (\ln \theta + 1)$.

Lemma 15 shows that $\Phi_\Delta(\pi)$ is upper bounded by $\bar{\Phi}(\pi)$. It is easy to show that $\bar{\Phi}(\pi)$ is decreasing in $\pi \geq 1$. Thus by setting $\bar{\Phi}(\bar{\pi}) = \Delta$, we can guarantee that CR-Pursuit($\bar{\pi}$) is feasible. Then we have the following result, which shows that the competitive ratio of CR-Pursuit($\bar{\pi}$) is $\ln \theta + \Omega(1)$.

Theorem 16. *Let $\bar{\pi} = (\ln \theta + 1)^2 / (\ln \theta + 3/4) < \ln \theta + 4/3$. The online algorithm CR-Pursuit($\bar{\pi}$) is feasible and is thus $\bar{\pi}$ -competitive.*

PROOF. With $\bar{\pi} = (\ln \theta + 1)^2 / (\ln \theta + 3/4)$, we have $\bar{\Phi}(\bar{\pi}) = \Delta$. From Lemma 15, we know that $\Phi_\Delta(\bar{\pi}) \leq \bar{\Phi}(\bar{\pi}) = \Delta$. Thus the results are immediate. \square

Note that $\bar{\pi} < \ln \theta + 4/3$, which is very close to the lower bound of $\ln \theta + 1$. It also improves beyond the result of $2(\ln \theta + 1)$ if we follow the characterization in Sec. 6; the improvement is because of the tighter bound through Lemma 14.

8 BEYOND THE WORST CASE MENTALITY

Our CR-Pursuit framework focuses only on achieving competitiveness under the worst case inputs. This may limit its applications as worst case inputs or situations may seldom occur in practice. Intuitively, a “better” online algorithm would sell more of its inventory when the incoming revenue function is “not adversarial”, i.e.,

being more opportunistic. By design, CR-Pursuit is pessimistic: it only maintains a fixed competitive ratio π^* for the whole trading period, even if some inputs are not adversarial. One way to improve the performance of CR-Pursuit for non-adversarial cases is as follows: instead of trying to keep the competitive ratio as π^* during the whole period, the online algorithm *adaptively* chooses a π_t to maintain at time t . This π_t is chosen as the smallest, yet attainable, competitive ratio at time t , given the previous inputs and outputs of the algorithm, and taking into account the possible inputs in future slots. This approach allows an online algorithm to instead pursue a competitive ratio more adaptive to the inputs, improving its average-case performance. We have recently applied the idea to develop online electric vehicle charging algorithms with optimal worst case and uniquely strong average-case performance [55].

To better illustrate this idea, consider the following example of one-way trading. Let the first price be $p(1) = M$. The original CR-Pursuit algorithm sells an $\Delta / (\ln \theta + 1)$ amount of inventory, and is satisfied with pursuing such competitive ratio at all time. The suggested algorithm in this section knows that the optimal offline value cannot increase any further, and would therefore sell all the inventory, i.e., we can set $\pi_1 = 1$. In this case, it will sell all the inventory in the first slot and achieve an offline-to-online revenue ratio of 1 for the particular input.

9 CONCLUDING REMARKS

Online optimization is an important line of research with wide ranging applications. It has been tackled by multiple algorithmic approaches over the previous decades, each proving successful for different problem variations, e.g., primal-dual approaches for online covering and packing problems or potential functions for the k -server problem.

In this work, we present a *novel* algorithmic framework for online optimization with inventory constraints. The framework “pursues” a bound on the competitive ratio, tracking the changes in the offline optimal algorithm and ensuring that the offline-to-online revenue ratio for the instance remains bounded throughout the entire period. This idea allows us to provide an nearly optimal algorithm for online optimization with inventory constraints as well as generalizations of the classical one-way trading problem. Specifically, our analysis and algorithms generalize naturally to one-way trading problems with price elasticity and concave revenue functions, yielding almost optimal (in terms of competitive ratio) online algorithms in those settings.

While our focus in this paper is on settings where inventory cannot be replenished, there is a wide range of applications with both selling periods and buying periods, like battery arbitrage in contingency markets. Usually in these markets, prices are highly affected by the selling quantities and also other factors that vary in time, which lead to unknown incoming revenue functions. We believe that the CR-Pursuit framework is promising for these problems as well, and can potentially be applicable to much broader classes of online optimization problems.

For example, our focus in this paper has been on worst case analysis but the CR-Pursuit framework can also be used to provide “beyond worst case” results by parameterizing the bound in different ways by, for example, utilizing properties of the instances relevant

to the application, and adaptively considering the input seen so far; see a recent example in [55]. Additionally, the framework can make use of randomization when pursuing the CR bound. This may allow improvement beyond the deterministic lower bound discussed in this paper, although it is an open question whether randomized algorithms can outperform deterministic algorithms for OOIC.

10 ACKNOWLEDGMENTS

We acknowledge the support received from Hong Kong University Grants Committee Theme-based Research Scheme Project No. T23-407/13-N and Collaborative Research Fund No. C7036-15G, NSF grant AST-134338, NSF AitF-1637598, NSF CNS-1518941, and NSF CPS-154471. In particular, John Pang wants to acknowledge the support from ASTAR, Singapore.

REFERENCES

- [1] Jacob D Abernethy, Elad Hazan, and Alexander Rakhlin. 2009. Competing in the dark: An efficient algorithm for bandit linear optimization. (2009).
- [2] H. Akhavan-Hejazi and H. Mohsenian-Rad. 2014. Optimal Operation of Independent Storage Systems in Energy and Reserve Markets With High Wind Penetration. *IEEE Transactions on Smart Grid* 5, 2 (March 2014), 1088–1097.
- [3] Susanne Albers. 2003. Online algorithms: a survey. *Mathematical Programming* 97, 1–2 (2003), 3–26.
- [4] Lachlan Andrew, Siddharth Barman, Katrina Ligett, Minghong Lin, Adam Meyerson, Alan Roytman, and Adam Wierman. 2013. A tale of two metrics: Simultaneous bounds on competitiveness and regret. In *Conference on Learning Theory*. 741–763.
- [5] Antonios Antoniadis, Neal Barcelo, Michael Nugent, Kirk Pruhs, Kevin Schewior, and Michele Scquizzato. 2016. Chasing convex bodies and functions. In *Latin American Symposium on Theoretical Informatics*. Springer, 68–81.
- [6] Sanjeev Arora, Elad Hazan, and Satyen Kale. 2012. The Multiplicative Weights Update Method: a Meta-Algorithm and Applications. *Theory of Computing* 8, 1 (2012), 121–164.
- [7] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. 1995. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *IEEE FOCS*. 322.
- [8] Y. Azar, N. Buchbinder, T. H. Chan, S. Chen, I. R. Cohen, A. Gupta, Z. Huang, N. Kang, V. Nagarajan, J. Naor, and D. Panigrahi. 2016. Online Algorithms for Covering and Packing Problems with Convex Objectives. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. 148–157.
- [9] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. 2008. Online auctions and generalized secretary problems. *ACM SIGecom Exchanges* 7, 2 (2008).
- [10] Nikhil Bansal, Martin Böhm, Marek Eliáš, Grigorios Koumoutsos, and Seeun William Umboh. 2018. Nested convex bodies are chaseable. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. 1253–1260.
- [11] Nikhil Bansal, Niv Buchbinder, and Joseph Seffi Naor. 2012. A primal-dual randomized algorithm for weighted paging. *Journal of the ACM (JACM)* 59, 4 (2012).
- [12] Nikhil Bansal, Anupam Gupta, Ravishankar Krishnaswamy, Kirk Pruhs, Kevin Schewior, and Cliff Stein. 2015. A 2-competitive algorithm for online convex optimization with switching costs. In *LIPICs-Leibniz International Proceedings in Informatics*, Vol. 40. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [13] H. Bogucka, M. Parzy, P. Marques, J. W. Mwangoka, and T. Forde. 2012. Secondary spectrum trading in TV white spaces. *IEEE Communications Magazine* 50, 11 (November 2012), 121–129.
- [14] Allan Borodin, Nathan Linial, and Michael E Saks. 1992. An optimal on-line algorithm for metrical task system. *Journal of the ACM (JACM)* 39, 4 (1992), 745–763.
- [15] Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning* 5, 1 (2012), 1–122.
- [16] Niv Buchbinder and Joseph Naor. 2005. Online primal-dual algorithms for covering and packing problems. In *European Symposium on Algorithms*. Springer, 689–701.
- [17] Niv Buchbinder, Joseph Seffi Naor, et al. 2009. The design of competitive online algorithms via a primal–dual approach. *Foundations and Trends® in Theoretical Computer Science* 3, 2–3 (2009), 93–263.
- [18] Niangjun Chen, Anish Agarwal, Adam Wierman, Siddharth Barman, and Lachlan LH Andrew. 2015. Online convex optimization using predictions. In *ACM SIGMETRICS Performance Evaluation Review*, Vol. 43. 191–204.
- [19] Francis YL Chin, Bin Fu, Jiuling Guo, Shuguang Han, Jueliang Hu, Minghui Jiang, Guohui Lin, Hing-Fung Ting, Luping Zhang, Yong Zhang, et al. 2015. Competitive algorithms for unbounded one-way trading. *Theoretical Computer Science* 607 (2015), 35–48.
- [20] YS Chow, Sigaiti Moriguti, Herbert Robbins, and SM Samuels. 1964. Optimal selection based on relative rank (the “secretary problem”). *Israel Journal of mathematics* 2, 2 (1964), 81–90.
- [21] Peter Damaschke, Phuong Hoai Ha, and Philippas Tsigas. 2009. Online search with time-varying price bounds. *Algorithmica* 55, 4 (2009), 619–642.
- [22] Nikhil R. Devanur and Kamal Jain. 2012. Online Matching with Concave Returns. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing (STOC '12)*. ACM, 137–144.
- [23] Ran El-Yaniv, Amos Fiat, Richard M Karp, and Gordon Turpin. 2001. Optimal search and one-way trading online algorithms. *Algorithmica* 30, 1 (2001), 101–139.
- [24] Moran Feldman and Rico Zenklusen. 2018. The submodular secretary problem goes linear. *SIAM J. Comput.* 47, 2 (2018), 330–366.
- [25] Amos Fiat. 1998. Online Algorithms: The State of the Art (Lecture Notes in Computer Science). (1998).
- [26] Amos Fiat and Manor Mendel. 2003. Better algorithms for unfair metrical task systems and applications. *SIAM J. Comput.* 32, 6 (2003), 1403–1422.
- [27] Amos Fiat, Yuval Rabani, and Yiftach Ravid. 1990. Competitive k-server algorithms. (1990).
- [28] Joel Friedman and Nathan Linial. 1993. On convex body chasing. *Discrete & Computational Geometry* 9, 3 (1993), 293–321.
- [29] Hiroshi Fujiwara, Kazuo Iwama, and Yoshiyuki Sekiguchi. 2011. Average-case competitive analyses for one-way trading. *Journal of Combinatorial Optimization* 21, 1 (2011), 83–107.
- [30] Edward F Grove. 1991. The harmonic online k-server algorithm is competitive. In *Proceedings of the twenty-third annual ACM symposium on Theory of computing*. 260–266.
- [31] Linqi Guo, John Pang, and Anwar Walid. [n. d.]. Dynamic service function chaining in SDN-enabled networks with middleboxes. In *2016 IEEE 24th International Conference on Network Protocols (ICNP)*. 1–10.
- [32] Linqi Guo, John Pang, and Anwar Walid. 2018. Joint Placement and Routing of Network Function Chains in Data Centers. In *IEEE INFOCOM*. 612–620.
- [33] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. 2007. Automated online mechanism design and prophet inequalities. In *AAAI*, Vol. 7. 58–65.
- [34] Elad Hazan, Amit Agarwal, and Satyen Kale. 2007. Logarithmic regret algorithms for online convex optimization. *Machine Learning* 69, 2–3 (2007), 169–192.
- [35] Sham M Kakade, Michael Kearns, Yishay Mansour, and Luis E Ortiz. 2004. Competitive algorithms for VWAP and limit order trading. In *Proceedings of ACM EC*. 189–198.
- [36] Anna R Karlin, Mark S Manasse, Larry Rudolph, and Daniel D Sleator. 1988. Competitive snoopy caching. *Algorithmica* 3, 1–4 (1988), 79–119.
- [37] Jiayin Li, Meikang Qiu, Zhong Ming, Gang Quan, Xiao Qin, and Zonghua Gu. 2012. Online optimization for scheduling preemptable tasks on IaaS cloud systems. *J. Parallel and Distrib. Comput.* 72, 5 (2012), 666–677.
- [38] Yingying Li, Guannan Qu, and Na Li. 2018. Using Predictions in Online Optimization with Switching Costs: A Fast Algorithm and A Fundamental Limit. In *IEEE Annual American Control Conference (ACC)*. 3008–3013.
- [39] Minghong Lin, Adam Wierman, Lachlan LH Andrew, and Eno Thereska. 2013. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Transactions on Networking (TON)* 21, 5 (2013), 1378–1391.
- [40] Minghong Lin, Adam Wierman, Alan Roytman, Adam Meyerson, and Lachlan LH Andrew. 2012. Online optimization with switching cost. *ACM SIGMETRICS Performance Evaluation Review* 40, 3 (2012), 98–100.
- [41] L.Lu, J.Tu, C.Chau, M.Chen, and X.Lin. 2013. Online Energy Generation Scheduling for Microgrids with Intermittent Energy Sources and Co-Generation. *Proceedings of ACM Sigmetrics* (2013).
- [42] Julian Lorenz, Konstantinos Panagiotou, and Angelika Steger. 2009. Optimal algorithms for k-search with application in option pricing. *Algorithmica* 55, 2 (2009), 311–328.
- [43] Tan Lu, Minghua Chen, and Lachlan LH Andrew. 2013. Simple and effective dynamic provisioning for power-proportional data centers. *IEEE Transactions on Parallel and Distributed Systems* 24, 6 (2013), 1161–1171.
- [44] Jacob Mattingley, Yang Wang, and Stephen Boyd. 2011. Receding horizon control. *IEEE Control Systems* 31, 3 (2011), 52–65.
- [45] Hanna Michalska and David Q Mayne. 1993. Robust receding horizon control of constrained nonlinear systems. *IEEE transactions on automatic control* 38, 11 (1993), 1623–1633.
- [46] John ZF Pang, Hu Fu, Won I Lee, and Adam Wierman. 2017. The Efficiency of Open Access in Platforms for Networked Cournot Markets. In *Proceedings of IEEE INFOCOM*.
- [47] John ZF Pang, Linqi Guo, and Steven H Low. 2017. Optimal load control for frequency regulation under limited control coverage. In *IREP Symposium*. 1–7.

- [48] L. Qian, F. Ye, L. Gao, X. Gan, T. Chu, X. Tian, X. Wang, and M. Guizani. 2011. Spectrum Trading in Cognitive Radio Networks: An Agent-Based Model under Demand Uncertainty. *IEEE Transactions on Communications* 59, 11 (November 2011), 3192–3203.
- [49] Xiaoqi Ren, Palma London, Juba Ziani, and Adam Wierman. 2018. Datum: Managing Data Purchasing and Data Placement in a Geo-Distributed Data Market. *IEEE/ACM Transactions on Networking* 26, 2 (2018), 893–905.
- [50] Aviad Rubinstein. 2016. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 324–332.
- [51] Karl H Schlag. 1998. Why imitate, and if so, how?: A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory* 78, 1 (1998), 130–156.
- [52] S. Shafiee, A. Bhuiya, A. U. Haque, H. Zareipour, and A. M. Knight. 2018. Economic Assessment of Energy Storage Systems in Alberta’s Energy and Operating Reserve Markets. In *2018 IEEE/PES Transmission and Distribution Conference and Exposition*. 1–5.
- [53] Shai Shalev-Shwartz and Sham M Kakade. 2009. Mind the duality gap: Logarithmic regret algorithms for online optimization. In *Advances in Neural Information Processing Systems*. 1457–1464.
- [54] Lin Yang, Mohammad H Hajiesmaili, Hanling Yi, and Minghua Chen. June 2017. Online Offering Strategies for Storage-Assisted Renewable Power Producer in Hour-Ahead Market. In *Proceedings of ACM SIGMETRICS*.
- [55] Hanling Yi*, Qiulin Lin*, and Minghua Chen. 2019. Balancing Cost and Dissatisfaction in Online EV Charging under Real-time Pricing. In *IEEE INFOCOM*. (The first two authors contribute equally to the work).
- [56] Wenming Zhang, Yinfeng Xu, Feifeng Zheng, and Yucheng Dong. 2012. Optimal algorithms for online time series search and one-way trading with interrelated prices. *Journal of combinatorial optimization* 23, 2 (2012), 159–166.
- [57] Ying Zhang, Mohammad H Hajiesmaili, Sinan Cai, Minghua Chen, and Qi Zhu. 2018. Peak-aware online economic dispatching for microgrids. *IEEE Transactions on Smart Grid* 9, 1 (2018), 323–335.

A APPENDIX

A.1 A Binary Search Algorithm for Computing λ^*

We summarize the algorithm in Algorithm 2.

Algorithm 2 A Binary search algorithm for Computing λ^*

```

1: if  $\max_{v_t \in V_t(0)} \sum_{t=1}^T v_t \leq \Delta$  then
2:   return  $\lambda^* = 0$ ;
3: else
4:   Pick  $\lambda_L = 0, \lambda_H = \max_{t \in \mathcal{T}} (g'_t(0))$ ;
5:   while  $|\lambda_L - \lambda_H| > \epsilon$  do
6:      $\lambda_M = \frac{\lambda_L + \lambda_H}{2}, v_t = 0, \forall t \in \mathcal{T}$ ;
7:     Compute  $\Delta_{\max} = \max_{v_t \in V_t(\lambda_M)} \sum_{t=1}^T v_t$ 
8:     Compute  $\Delta_{\min} = \min_{v_t \in V_t(\lambda_M)} \sum_{t=1}^T v_t$ .
9:     if  $\Delta_{\min} > \Delta$  then
10:       $\lambda_L = \lambda_M$ ;
11:     end if
12:     if  $\Delta_{\max} < \Delta$  then
13:       $\lambda_H = \lambda_M$ ;
14:     end if
15:     if  $\Delta_{\min} \leq \Delta \leq \Delta_{\max}$  then
16:      break;
17:     end if
18:   end while
19:   return  $\lambda^* = \lambda_M$ ;
20: end if

```

A.2 Proof of Proposition 1

PROOF. We prove this theorem by investigating the KKT conditions of problem OOIC and exploring the structure of the optimal solution.

The Lagrangian for problem OOIC is defined as

$$L(v, \lambda, \mu) = \sum_{t=1}^T g_t(v_t) + \lambda \left(\Delta - \sum_{t=1}^T v_t \right) + \sum_{t=1}^T v_t \mu(t),$$

where $\lambda \geq 0$ and $\mu(t) \geq 0, \forall t \in [T]$ are the Lagrangian multipliers. The following KKT conditions give us a set of necessary and sufficient conditions for optimality:

$$g'_t(v_t) - \lambda + \mu(t) = 0, \quad \forall t \in [T],$$

$$\sum_{t=1}^T v_t \leq \Delta,$$

$$v_t \geq 0, \quad \forall t \in [T],$$

$$\mu(t) \geq 0, \quad \forall t \in [T],$$

$$\lambda \geq 0,$$

$$v_t \mu(t) = 0, \quad \forall t \in [T],$$

$$\lambda \left(\sum_{t=1}^T v_t - \Delta \right) = 0.$$

Suppose v^*, μ^* and λ^* are the optimal solutions that satisfy the KKT conditions. Denote the set $\mathcal{T}_0 = \{t | v_t^* > 0, \forall t \in [T]\}$, then according to the KKT conditions, we have

$$\mu^*(t) = 0, \quad \forall t \in \mathcal{T}_0, \quad (19)$$

$$\lambda^* \left(\sum_{t \in \mathcal{T}_0} v_t^* - \Delta \right) = 0, \quad (20)$$

$$g'_t(v_t^*) - \lambda^* = 0, \quad \forall t \in \mathcal{T}_0, \quad (21)$$

Since g'_t is concave, $g'_t(\cdot)$ is non-increasing in v_t . According to (21) we have

$$g'_t(0) \geq g'_t(v_t^*) = \lambda^*, \quad \forall t \in \mathcal{T}_0;$$

namely,

$$g'_t(0) \geq \lambda^* \quad \forall t \in \mathcal{T}_0. \quad (22)$$

Thus given a λ^* , we can use (22) to determine the set \mathcal{T}_0 .

For ease of presentation, we denote

$$V_t(\lambda) = \{v | g'_t(v) = \lambda, v \in [0, \Delta]\}.$$

Now consider the following two cases:

(1) $\Delta \geq \max_{v_t \in V_t(0)} \sum_{t=1}^T v_t$. In this case, we observe that the solution

$$v_t^* \in V_t(0), \forall t \in [T],$$

$$\lambda^* = 0,$$

$$\mu^*(t) = 0, \forall t \in [T],$$

satisfies the KKT conditions, thus it is the optimal solution.

(2) $\Delta < \max_{v_t \in V_t(0)} \sum_{t=1}^T v_t$. In this case, we must have $\lambda^* > 0$. According to (20) and (21), we have

$$v_t^* \in V_t(\lambda^*) \text{ and } \sum_{t=1}^T v_t^* = \Delta.$$

It is straightforward to check that $v_t, \forall t \in \mathcal{T}_0$, is non-increasing w.r.t. λ . Meanwhile, according to (22), we know that the size of set \mathcal{T}_0 is non-increasing w.r.t. λ . Putting together these two observations, we conclude that $\sum_{t \in \mathcal{T}_0} v_t$ is non-increasing w.r.t. λ . Thus given $\Delta > 0$, there exists a unique $\lambda = \lambda^*$ that satisfies $\sum_{t \in \mathcal{T}_0} v_t^* = \Delta$. Since KKT conditions are necessary and sufficient for optimality of convex problems, we can conclude that λ^* is the optimal dual solution. \square

A.3 Proof of Lemma 2

PROOF OF LEMMA 2. We prove this lemma in the following two steps:

Step I, we prove that $\eta_{OPT}(\sigma^{[1:t]}) - \eta_{OPT}(\sigma^{[1:t-1]}) \geq g_t(\tilde{v}_t) - \lambda_t \tilde{v}_t$. To see this, we denote optimal solution at time $\tau \in [t]$ under input $\sigma^{[1:t]}$ as \tilde{v}_τ . Note that $\tilde{v}_\tau \in V_\tau(\lambda_t), \tau \in [t]$ or $\tilde{v}_\tau = 0$ if $V_\tau(\lambda_t) = \emptyset$. Similarly, denote optimal solution at time $\tau \in [t-1]$ under input $\sigma^{[1:t-1]}$ as \bar{v}_τ . Note that $\bar{v}_\tau \in V_\tau(\lambda_{t-1}), \tau \in [t-1]$ or $\bar{v}_\tau = 0$ if $V_\tau(\lambda_{t-1}) = \emptyset$. Also $\tilde{v}_\tau \leq \bar{v}_\tau, \tau \in [t-1]$ (by the non-increasing of $g'_t(v)$ and $\lambda_t \geq \lambda_{t-1}$). Then we have

$$\begin{aligned} \eta_{OPT}(\sigma^{[1:t]}) - \eta_{OPT}(\sigma^{[1:t-1]}) &= \sum_{\tau=1}^t g_\tau(\tilde{v}_\tau) - \sum_{\tau=1}^{t-1} g_\tau(\bar{v}_\tau) \\ &= g_t(\tilde{v}_t) + \sum_{\tau=1}^{t-1} (g_\tau(\tilde{v}_\tau) - g_\tau(\bar{v}_\tau)) \\ &\stackrel{(a)}{\geq} g_t(\tilde{v}_t) + \sum_{\tau=1}^{t-1} \lambda_t (\tilde{v}_\tau - \bar{v}_\tau) \\ &\stackrel{(b)}{\geq} g_t(\tilde{v}_t) - \lambda_t \tilde{v}_t. \end{aligned}$$

For (a), it comes from the concavity of $g_\tau(v)$ and $\tilde{v}_\tau \leq \bar{v}_\tau, \tau \in [t-1]$. For (b), we claim that $\sum_{\tau=1}^{t-1} \bar{v}_\tau \leq \sum_{\tau=1}^t \tilde{v}_\tau$. To see this, when $\lambda_t = 0$, we must have $\lambda_{t-1} = 0$. In this case, $\tilde{v}_\tau = \bar{v}_\tau, \forall \tau \in [t-1]$ and thus we have $\sum_{\tau=1}^{t-1} \bar{v}_\tau \leq \sum_{\tau=1}^t \tilde{v}_\tau$. When $\lambda_t > 0$, from the KKT conditions in (20), we have $\sum_{\tau=1}^t \tilde{v}_\tau = \Delta \geq \sum_{\tau=1}^{t-1} \bar{v}_\tau$. Then we conclude that $\sum_{\tau=1}^{t-1} \bar{v}_\tau \leq \sum_{\tau=1}^t \tilde{v}_\tau$ and consequently, we have $\sum_{\tau=1}^{t-1} (\tilde{v}_\tau - \bar{v}_\tau) \geq -\tilde{v}_t$.

Step II, we prove that $\eta_{OPT}(\sigma^{[1:t]}) - \eta_{OPT}(\sigma^{[1:t-1]}) \leq g_t(\tilde{v}_t) - \lambda_{t-1} \tilde{v}_t \leq g_t(\tilde{v}_t) \leq g_t(\hat{v}_t)$. Similarly, we have

$$\begin{aligned} \eta_{OPT}(\sigma^{[1:t]}) - \eta_{OPT}(\sigma^{[1:t-1]}) &= g_t(\tilde{v}_t) + \sum_{\tau=1}^{t-1} (g_\tau(\tilde{v}_\tau) - g_\tau(\bar{v}_\tau)) \\ &\stackrel{(a)}{\leq} g_t(\tilde{v}_t) + \sum_{\tau=1}^{t-1} \lambda_{t-1} (\tilde{v}_\tau - \bar{v}_\tau) \\ &\stackrel{(b)}{=} g_t(\tilde{v}_t) - \lambda_{t-1} \tilde{v}_t \\ &\leq g_t(\tilde{v}_t) \leq g_t(\hat{v}_t). \end{aligned}$$

For (a), it is by the concavity of $g_\tau: g_\tau(\tilde{v}_\tau) \leq g_\tau(\bar{v}_\tau) + \lambda_{t-1}(\tilde{v}_\tau - \bar{v}_\tau)$ (Note that $\lambda_\tau = g'_\tau(\bar{v}_\tau)$) and $\lambda_{t-1} \geq \lambda_\tau, \forall \tau \in [t-1]$. For (b), when $\lambda_{t-1} = 0$, it holds immediately; when $\lambda_{t-1} > 0$, we have $\sum_{\tau=1}^t \tilde{v}_\tau = \Delta = \sum_{\tau=1}^{t-1} \bar{v}_\tau$, which implies $\sum_{\tau=1}^{t-1} (\tilde{v}_\tau - \bar{v}_\tau) = -\tilde{v}_t$. \square

A.4 Proof of Lemma 3

PROOF. Denote the input under $\tilde{\sigma}$ as g_t . Denote the input under $\bar{\sigma}$ as \bar{g}_t . The optimal dual variable under $\tilde{\sigma}^{[1:t]}$ (resp. $\bar{\sigma}^{[1:t]}$) as λ_t (resp. $\bar{\lambda}_t$). We have,

$$g_t = \bar{g}_t, \forall t \leq \tau - 1 \vee t \geq \tau + 2.$$

Besides, $g_\tau = \bar{g}_{\tau+1}, g_{\tau+1} = \bar{g}_\tau$. Let v_t (resp. \bar{v}_t) be the optimal offline solution at time t given the input $\tilde{\sigma}^{[1:t]}$ (resp. $\bar{\sigma}^{[1:t]}$).

1) If $\lambda_\tau \leq \bar{\lambda}_\tau$, then

$$\begin{aligned} \eta_{OPT}(\tilde{\sigma}^{[1:\tau]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau-1]}) &\stackrel{(a)}{\geq} g_\tau(v_\tau) - \lambda_\tau v_\tau \\ &\stackrel{(b)}{\geq} g_\tau(\bar{v}_{\tau+1}) - \lambda_\tau \bar{v}_{\tau+1} \\ &\stackrel{(c)}{\geq} g_\tau(\bar{v}_{\tau+1}) - \bar{\lambda}_\tau \bar{v}_{\tau+1} \\ &\stackrel{(a)}{\geq} \eta_{OPT}(\tilde{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau]}). \end{aligned}$$

For (a), it is by lemma 2. For (b), it is by the concavity of g_t and for (c), it is by $\lambda_\tau \leq \bar{\lambda}_\tau$.

2) If $\lambda_\tau \geq \bar{\lambda}_\tau$, then similarly

$$\begin{aligned} \eta_{OPT}(\tilde{\sigma}^{[1:\tau]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau-1]}) &\stackrel{(a)}{\geq} g_{\tau+1}(\bar{v}_\tau) - \bar{\lambda}_\tau \bar{v}_\tau \\ &\stackrel{(b)}{\geq} g_{\tau+1}(\bar{v}_\tau) - \lambda_\tau \bar{v}_\tau \\ &\stackrel{(c)}{\geq} g_{\tau+1}(v_{\tau+1}) - \lambda_\tau v_{\tau+1} \\ &\stackrel{(a)}{\geq} \eta_{OPT}(\tilde{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau]}). \end{aligned}$$

For (a), it is by lemma 2. For (b), it is by $\lambda_\tau \geq \bar{\lambda}_\tau$. For (c), it is by the concavity of g_t . Also, with

$$\begin{aligned} &\eta_{OPT}(\tilde{\sigma}^{[1:\tau]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau-1]}) + \eta_{OPT}(\tilde{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau]}) \\ &= \eta_{OPT}(\tilde{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau-1]}) \\ &= \eta_{OPT}(\tilde{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau-1]}) \\ &= \eta_{OPT}(\tilde{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau]}) + \eta_{OPT}(\tilde{\sigma}^{[1:\tau]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau-1]}), \end{aligned}$$

we can have

$$\eta_{OPT}(\tilde{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau]}) \leq \eta_{OPT}(\tilde{\sigma}^{[1:\tau]}) - \eta_{OPT}(\tilde{\sigma}^{[1:\tau-1]}). \quad \square$$

A.5 Proof of Lemma 7

PROOF OF LEMMA 7. Suppose an arbitrary $\tilde{\sigma} \in \arg \max_{\sigma} \sum_t v_t$, under which $g'_t(v_t)$ is not non-decreasing in t , where v_t is the selling quantity of CR-Pursuit(π) under $\tilde{\sigma}$. That is, exist a $\tau, g'_\tau(v_\tau) > g'_{\tau+1}(v_{\tau+1})$. Denote the optimal dual variables under $\tilde{\sigma}^{[1:t]}$ as λ_t . Note that λ_t is non-decreasing in t . Without loss of generality, we assume that $\lambda_t < \lambda_{t+1}$ or $\lambda_t = \lambda_{t+1} = 0, \forall t$. We construct a new input sequence $\bar{\sigma}$ by interchanging g_τ and $g_{\tau+1}$ in $\tilde{\sigma}$ and denote the input under $\bar{\sigma}$ as \bar{g}_t , the output of CR-Pursuit(π^*) under $\bar{\sigma}$ as \bar{v}_t . The optimal dual variable under $\bar{\sigma}^{[1:t]}$ as $\bar{\lambda}_t$. By definition, we can easily observe that,

$$\eta_{OPT}(\tilde{\sigma}^{[1:t]}) = \eta_{OPT}(\bar{\sigma}^{[1:t]}), \forall t \leq \tau - 1 \vee t \geq \tau + 1;$$

$$v_t = \bar{v}_t, \forall t \leq \tau - 1 \vee t \geq \tau + 2;$$

$$g_t = \bar{g}_t, \forall t \leq \tau - 1 \vee t \geq \tau + 2.$$

Besides, $g_\tau = \bar{g}_{\tau+1}, g_{\tau+1} = \bar{g}_\tau$. We claim that $\bar{\sigma} \in \arg \max_{\sigma} \sum_t v_t$ and $\bar{g}'_\tau(\bar{v}_\tau) = g'_{\tau+1}(v_{\tau+1}) < g'_\tau(v_\tau) = \bar{g}'_{\tau+1}(\bar{v}_{\tau+1})$. To see this, consider the following two cases.

(1) $\lambda_\tau = \lambda_{\tau+1} = 0$. Under this case, we have

$$\begin{aligned} \eta_{OPT}(\bar{\sigma}^{[1:\tau]}) - \eta_{OPT}(\bar{\sigma}^{[1:\tau-1]}) &= \eta_{OPT}(\bar{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\bar{\sigma}^{[1:\tau]}) \\ &= g_\tau(\hat{v}_\tau), \end{aligned}$$

where $\hat{v}_\tau = \arg \max_v g_\tau(v)$. Then $v_\tau = \bar{v}_{\tau+1}$. Similarly, we have $v_{\tau+1} = \bar{v}_\tau$. $\sum_t v_t = \sum_t \bar{v}_t$. We conclude that $\bar{\sigma} \in \arg \max_{\sigma} \sum_t v_t$ and $\bar{g}'_\tau(\bar{v}_\tau) = g'_{\tau+1}(v_{\tau+1}) < g'_\tau(v_\tau) = \bar{g}'_{\tau+1}(\bar{v}_{\tau+1})$.

(2) $0 \leq \lambda_\tau < \lambda_{\tau+1}$. First, we have

$$\begin{aligned} g_\tau(v_\tau) + g_{\tau+1}(v_{\tau+1}) &= \frac{\eta_{OPT}(\bar{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\bar{\sigma}^{[1:\tau-1]})}{\pi^*} \\ &= \frac{\eta_{OPT}(\bar{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\bar{\sigma}^{[1:\tau-1]})}{\pi^*} \\ &= g_{\tau+1}(\bar{v}_\tau) + g_\tau(\bar{v}_{\tau+1}), \end{aligned}$$

which implies

$$g_\tau(v_\tau) - g_\tau(\bar{v}_{\tau+1}) = g_{\tau+1}(\bar{v}_\tau) - g_{\tau+1}(v_{\tau+1}).$$

Second, we claim that $\bar{v}_{\tau+1} \leq v_\tau$. From Lemma 3, we have

$$\eta_{OPT}(\bar{\sigma}^{[1:\tau+1]}) - \eta_{OPT}(\bar{\sigma}^{[1:\tau]}) \leq \eta_{OPT}(\bar{\sigma}^{[1:\tau]}) - \eta_{OPT}(\bar{\sigma}^{[1:\tau-1]}).$$

Then $g_\tau(v_\tau) \geq g_\tau(\bar{v}_{\tau+1})$ and $\bar{v}_{\tau+1} \leq v_\tau$ are straightforward.

Third, we show $g_\tau(v_\tau) = g_\tau(\bar{v}_{\tau+1})$ and thus $\bar{v}_{\tau+1} = v_\tau$ by contradiction. Suppose $g_\tau(v_\tau) > g_\tau(\bar{v}_{\tau+1})$ and thus $\bar{v}_{\tau+1} < v_\tau$. We show that $\sum_t v_t < \sum_t \bar{v}_t$ which contradict the fact that $\bar{\sigma} \in \arg \max_{\sigma} \sum_t v_t$. To see this, observe that we have

$$\begin{aligned} g'_{\tau+1}(v_{\tau+1})(v_\tau - \bar{v}_{\tau+1}) &\stackrel{(a)}{<} -g'_\tau(v_\tau)(\bar{v}_{\tau+1} - v_\tau) \\ &\stackrel{(b)}{\leq} g_\tau(v_\tau) - g_\tau(\bar{v}_{\tau+1}) \\ &= g_{\tau+1}(\bar{v}_\tau) - g_{\tau+1}(v_{\tau+1}) \\ &\stackrel{(b)}{\leq} g'_{\tau+1}(v_{\tau+1})(\bar{v}_\tau - v_{\tau+1}). \end{aligned}$$

For (a), it is by $g'_\tau(v_\tau) > g'_{\tau+1}(v_{\tau+1}) \geq \lambda_{\tau+1} > 0$ and $\bar{v}_{\tau+1} < v_\tau$. For (b), it is from the concavity of g_τ . As $g'_{\tau+1}(v_{\tau+1}) \geq \lambda_{\tau+1} > 0$, we have

$$v_\tau + v_{\tau+1} < \bar{v}_\tau + \bar{v}_{\tau+1},$$

which leads to $\sum_t v_t < \sum_t \bar{v}_t$.

So we conclude that $g_\tau(v_\tau) = g_\tau(\bar{v}_{\tau+1})$ and thus $\bar{v}_{\tau+1} = v_\tau$. Consequently, $g_\tau(v_{\tau+1}) = g_{\tau+1}(\bar{v}_\tau)$ and thus $\bar{v}_\tau = v_{\tau+1}$. It is then straightforward that

$$\bar{\sigma} \in \arg \max_{\sigma} \sum_t v_t,$$

and

$$\bar{g}'_\tau(\bar{v}_\tau) = g'_{\tau+1}(v_{\tau+1}) < g'_\tau(v_\tau) = \bar{g}'_{\tau+1}(\bar{v}_{\tau+1}).$$

By continuously interchanging g_τ and $g_{\tau+1}$ that fails to satisfy $g'_{\tau+1}(v_\tau) \leq g'_\tau(v_{\tau+1})$, we finally attain a sequence in $\arg \max_{\sigma} \sum_t v_t$ such that $g'_t(v_t)$ is non-decreasing in t . \square

A.6 Proof of Lemma 9

PROOF. First, from Lemma 2, we easily conclude that $\bar{v}_t \leq \hat{v}_t$, where \hat{v}_t is the optimizer of $g_t(\cdot)$. By the concavity of $g_t(\cdot)$, we have

$$g_t(\bar{v}_t) \geq \frac{\bar{v}_t}{\hat{v}_t} g_t(\hat{v}_t) + \left(1 - \frac{\bar{v}_t}{\hat{v}_t}\right) g_t(0) \geq \frac{\bar{v}_t}{\hat{v}_t} g_t(\hat{v}_t),$$

which then gives $\bar{v}_t \leq \frac{g_t(\hat{v}_t)}{g_t(\bar{v}_t)/\hat{v}_t}$. Then, using the definition of c , we arrive at

$$\bar{v}_t \leq \frac{p(t)}{g_t(\hat{v}_t)/\hat{v}_t} \frac{g_t(\bar{v}_t)}{p(t)} \leq c \frac{g_t(\bar{v}_t)}{p(t)}.$$

\square

A.7 Proof of Lemma 10

PROOF. For ease of presentation, define

$$x_p \triangleq \frac{\pi}{\Delta} \sum_{\{t: p(t) \leq p\}} g_t(\bar{v}_t).$$

It is then equivalent to show that $x_p \leq p$. Define $T_1 \triangleq \min\{t : p(t) > p, \forall \tau \geq t\} - 1$, i.e., for any $t > T_1$, we have $p(t) > p$, or equivalently if $p(t) \leq p$, then $t \leq T_1$. By definition, x_p is determined by $\sigma^{[1:T_1]}$ only. Thus, in this proof, we only focus on the input horizon $t \in [T_1]$.

We first consider a special case when $p(t) \leq p, \forall t \in [T_1]$. By that $g_t(v), \forall t \in [T_1]$ are concave functions, we have

$$\eta_{OPT}(\sigma^{[1:T_1]}) = \sum_{t=1}^{T_1} g_t(v_t^*) \leq \sum_{t=1}^{T_1} (g_t(0) + g'_t(0)v_t^*) = \sum_{t=1}^{T_1} p(t)v_t^*,$$

where $v_t^*, t \in [T_1]$ are the solution of the optimal offline algorithm under input $\sigma^{[1:T_1]}$. Then according to (9) and that $\eta_{OPT}(\sigma^{[1:T_1]}) \leq p \cdot \Delta$, we have

$$x_p = \frac{\pi}{\Delta} \sum_{t=1}^{T_1} g_t(\bar{v}_t) = \frac{1}{\Delta} \eta_{OPT}(\sigma^{[1:T_1]}) \leq p.$$

We now consider the general cases, where there could be some slot(s) $\tau \in [T_1]$ such that $p(\tau) > p$. Then we construct a new input sequence $\bar{\sigma}$ by interchange g_τ and $g_{\tau+1}$ in σ . Denote the input under $\bar{\sigma}$ as \bar{g}_t . Let $\bar{x}_p, \bar{p}(t)$ be the corresponding variables under $\bar{\sigma}$.

To show that $x_p \leq p$, we first show $x_p \leq \bar{x}_p$. By definition, we observe that,

$$\eta_{OPT}(\sigma^t) = \eta_{OPT}(\bar{\sigma}^t), \forall t \leq \tau - 1 \vee t \geq \tau + 1;$$

$$g_t = \bar{g}_t, \forall t \leq \tau - 1 \vee t \geq \tau + 2.$$

Besides, $g_\tau = \bar{g}_{\tau+1}, g_{\tau+1} = \bar{g}_\tau$. We discuss two cases.

- When $p(\tau + 1) > p$: it is easy to see that $\bar{x}_p = x_p$.

We then prove $x_p \leq p$ as follows: We continuously interchange with $p(\tau) > p$ with the input at its next slot until all the slots with $p(t) \leq p$ is at the front of it. At the meantime, x_p keeps on non-decreasing. Finally, we get a σ' , in which the price at each slot in $[T_1']$ (T_1' is corresponding to T_1 but defined under σ') is less or equal to p , and $x_p \leq x_p'$. Since in $\sigma', p \geq p(t), \forall t$, from our analysis in the first part (special case), we have $x_p' \leq p$. It then follows that $x_p \leq p$.

- When $p(\tau + 1) \leq p$: we have

$$x_p - \bar{x}_p = \frac{\eta_{OPT}(\sigma^{[1:\tau+1]}) - \eta_{OPT}(\sigma^{[1:\tau]})}{\Delta} - \frac{\eta_{OPT}(\bar{\sigma}^{[1:\tau]}) - \eta_{OPT}(\bar{\sigma}^{[1:\tau-1]})}{\Delta}$$

(a)
 ≤ 0 ,

where the step (a) is because of Lemma 3.

Next, we prove $x_p \leq p$. We continuously interchange with $p(\tau) > p$ with the input at its next slot until all the slots with $p(t) \leq p$ is at the front of it. At the meantime, x_p keeps on non-decreasing. Finally, we get a σ' , in which the price at each slot in $[T_1']$ (T_1' is corresponding to T_1 but defined under σ') is less or equal to p , and $x_p \leq x_p'$. Since in σ' , $p \geq p(t)$, $\forall t$, from our analysis in the first part (special case), we have $x_p' \leq p$. It then follows that $x_p \leq p$. \square

A.8 Proof of Lemma 11

PROOF. Suppose in $\sigma^{[1:T]}$, $p(t)$ takes n different values, which are denoted as $m \leq p_1 \leq p_2 \leq \dots \leq p_n \leq M$. And define $y_i \triangleq \sum_{t, p(t)=p_i} \frac{\pi}{\Delta} g_t(\bar{v}_t)$. Note that we have

$$\sum_{t=1}^T \frac{g_t(\bar{v}_t)}{p(t)} = \frac{\Delta}{\pi} \sum_{i=1}^n \frac{y_i}{p_i}.$$

From Lemma 10, we have $\sum_{j=1}^i y_j = x_{p_i} \leq p_i$.

Consider the following optimization problem:

$$\begin{aligned} \max \quad & \sum_{i=1}^n \frac{y_i}{p_i} \\ \text{s.t.} \quad & \sum_{j=1}^i y_j \leq p_i, i \in [n] \\ & y_i \geq 0, i \in [n]. \end{aligned}$$

The KKT conditions are sufficient and necessary conditions for optimality for the above convex problem. Denote $\mu_i \geq 0, i \in [n]$ as the dual variables, then the KKT conditions can be expressed as:

$$\frac{1}{p_i} - \sum_{j=1}^{n+1-i} \mu_j = 0, \forall i \in [n], \quad (23)$$

$$\mu_i(p_i - \sum_{j=1}^i y_j) = 0, \forall i \in [n], \quad (24)$$

$$\mu_i \geq 0, \forall i \in [n],$$

$$y_i \geq 0, \forall i \in [n].$$

From (23), we know that $\mu_i > 0$ for all $i \in [n]$. Thus from (24), we have

$$p_i - \sum_{j=1}^i y_j = 0, \forall i \in [n].$$

Thus we know the optimal primal solution is

$$y_i = p_i - p_{i-1}, \forall i \in [n],$$

where $p_0 = 0$. And the optimal objective value equals to $\sum_{i=1}^n \frac{p_i - p_{i-1}}{p_i}$. So

$$\begin{aligned} \sum_{t=1}^T \frac{g_t(\bar{v}_t)}{p(t)} &= \frac{\Delta}{\pi} \sum_{i=1}^n \frac{y_i}{p_i} \\ &\leq \frac{\Delta}{\pi} \sum_{i=1}^n \frac{p_i - p_{i-1}}{p_i} \\ &= \frac{\Delta}{\pi} \left(\frac{p_1}{p_1} + \sum_{i=2}^n \frac{p_i - p_{i-1}}{p_i} \right) \\ &\leq \frac{\Delta}{\pi} \left(1 + \int_{p_1}^{p_n} \frac{1}{x} dx \right) \\ &\leq \frac{\Delta}{\pi} (1 + \ln \theta). \end{aligned}$$

This completes our proof. \square

A.9 Proof of Lemma 12

PROOF. We show that any input $\sigma^{[1:T]}$ is equivalent to (in the sense that the behaviors of both offline algorithm and the proposed online algorithm remain unchanged) an increasing price sequence as the following:

$$m \leq p_1 < p_2 < \dots < p_n \leq M, \quad (25)$$

where $n \leq T$. According to (16), CR-Pursuit(π) will sell only when the current price is larger than the highest price in history. Thus for any input $\sigma^{[1:T]}$, we can delete the slots when CR-Pursuit(π) does not sell, and the outputs of CR-Pursuit(π) is then equivalent to the resulting increasing price sequence. \square

A.10 Proof of Lemma 14

PROOF. By Lemma 2 and definition of $g_t(\bar{v}_t)$, we know that $\pi g_t(\bar{v}_t) = \eta_{OPT}(\sigma^{[1:t]}) - \eta_{OPT}(\sigma^{[1:t-1]}) \leq g_t(\hat{v}_t)$ and then $\bar{v}_t \leq \hat{v}_t$, where \hat{v}_t as the optimizer of $g_t(v_t)$. To simplify the explanation, let $k = g_t(\bar{v}_t)$ and $\alpha = \frac{f(\bar{v}_t)}{\bar{v}_t}$.

Define $\tilde{g}_t(v_t) = (p(t) - \alpha v_t) v_t$. By the convexity of $f_t(\cdot)$, $f_t(0) = 0$, we have $f_t(v_t) \leq \alpha v_t, \forall v_t \leq \hat{v}_t$. Then $g_t(v_t) \geq \tilde{g}_t(v_t), \forall v_t \leq \hat{v}_t$.

Suppose \bar{v}_t is the smaller solution satisfying $\tilde{g}_t(\bar{v}_t) = k$, i.e.,

$$\bar{v}_t = \frac{p(t) - \sqrt{p^2(t) - 4\alpha k}}{2\alpha} = \frac{2k}{p(t) \left(1 + \sqrt{1 - \frac{4\alpha k}{p^2(t)}} \right)}.$$

By observing

$$k\pi \leq g_t(\hat{v}_t) = \tilde{g}_t(\hat{v}_t) \leq \frac{p^2(t)}{4\alpha},$$

we have $\frac{4\alpha k}{p^2} \leq \frac{1}{\pi}$ (note that this also implies the existence of \bar{v}_t). We then easily conclude

$$\bar{v}_t \leq \frac{2k}{p(t) \left(1 + \sqrt{1 - \frac{1}{\pi}} \right)}.$$

We claim $\bar{v}_t \leq \hat{v}_t$. If $\bar{v}_t > \hat{v}_t$, we have $\bar{v}_t \leq \hat{v}_t < \bar{v}_t$; Otherwise, $\bar{v}_t \leq \hat{v}_t$, we have

$$k = g_t(\bar{v}_t) = \tilde{g}_t(\bar{v}_t) \leq g_t(\hat{v}_t).$$

Following $g_t(v_t)$ is increasing in $[0, \hat{v}_t]$, we conclude $\bar{v}_t \leq \hat{v}_t$.

Finally, we conclude

$$\bar{v}_t \leq \frac{2k}{p(t) \left(1 + \sqrt{1 - \frac{1}{\pi}}\right)} = \frac{2}{\left(1 + \sqrt{1 - \frac{1}{\pi}}\right)} \frac{g_t(\bar{v}_t)}{p(t)}.$$

□

A.11 Proof of Lemma 15

PROOF. From Lemma 14, we have

$$\Phi_\Delta(\pi) = \max_{\sigma^{[1:T]}} \sum_{t=1}^T v_t \leq \frac{2}{\left(1 + \sqrt{1 - \frac{1}{\pi}}\right)} \sum_{t=1}^T \frac{g_t(\bar{v}_t)}{p(t)}.$$

By Lemma 11, we know that

$$\sum_{t=1}^T \frac{g_t(\bar{v}_t)}{p(t)} \leq \frac{\Delta}{\pi} (1 + \ln \theta). \quad (26)$$

Then we can bound $\Phi_\Delta(\pi)$ as

$$\begin{aligned} \Phi_\Delta(\pi) &\leq \frac{2}{\left(1 + \sqrt{1 - \frac{1}{\pi}}\right)} \sum_{t=1}^T \frac{g_t(\bar{v}_t)}{p(t)} \\ &\leq \frac{2\Delta}{\pi \left(1 + \sqrt{1 - \frac{1}{\pi}}\right)} (1 + \ln \theta). \end{aligned}$$

This completes our proof.

□