European Journal of Operational Research 252 (2016) 879–887

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Competitive difference analysis of the one-way trading problem with limited information

Wei Wang[a], Liying Wang[b], Yingjie Lan[c,*], Jean X. Zhang[d]

[a] Business School, Nankai University, Tianjin 300071, China
[b] College of Business Administration, Capital University of Economics and Business, Beijing 100070, China
[c] HSBC Business School, Peking University, University Town, Shenzhen 518055, China
[d] School of Business, Virginia Commonwealth University, Richmond, Virginia 23284, USA

ARTICLE INFO

Article history:
Received 22 July 2014
Accepted 19 January 2016
Available online 27 January 2016

Keywords:
Robustness and sensitivity analysis
One-way trading
Competitive analysis
Online algorithms

ABSTRACT

We consider robust one-way trading with limited information on price fluctuations. Our analysis finds the best guarantee of difference from the optimal offline performance. We provide closed-form solution, and reveal for the first time all possible worst-case scenarios. Numerical experiments show that our policy is more tolerant of information inaccuracy than Bayesian policies, and can earn higher average revenue than other robust policies while keeping a lower standard deviation.

1. Introduction

The one-way trading problem is well studied in the literature, where one-way trading simply means either selling or buying, but not both. A certain quantity of goods must be traded within a given time horizon during which the price fluctuates. The trader has to decide on the amount to trade at each time point as the current price becomes available, aiming to optimize the total trading value, that is, either to maximize the selling revenue or to minimize the buying cost. Typical applications of one-way trading include inventory procurement, merchandising, stock trading, currency exchange, online auctions, and so on. Such a trading process is an online decision-making process, in which irrevocable decisions must be made on the spot at each time point of action, in the face of future uncertainties. The precise rule on how to make the trading decisions during the process is referred to as a trading policy or algorithm. Clearly, how much is known about the future prices is crucial for the trader. When the exact path of future price fluctuation is known in advance, it is equivalent to the offline decision-making that happens in the end of the time horizon after all information is known. The optimal offline policy is simple: to sell all at the highest price or to buy all at the lowest price. The problem only becomes interesting when less is known about future prices, and different problem formulations are proposed in the literature with various degrees of information availability on future prices.

Earlier studies assumed probabilistic distribution on price uncertainty, and formulated stochastic optimization problems to design trading policies based on Bayes' rule (Lippman & McCall, 1976, 1981). The distribution of future prices can be estimated from historical data. However, this approach has some limitations due to difficulties in estimating the distribution, besides its risk neutrality assumption. Firstly, there are ample cases where appropriate historical data is unavailable, such as with newly launched products. Secondly, even if there is historical data, it may not be used to predict the distribution of future prices, when the price fluctuation is following some non-stationary random processes or due to some emergency situation which impacts future prices structurally. Finally, there are different ways to estimate the distribution from data and the particular choice of estimation could mislead the trader. Thus when there are reasons to doubt the soundness of the estimated distribution, it is valuable for the trader to find a robust trading policy based on less but more reliable information about future prices, which gives occasion to the research on trading policies with only limited price information.

A formulation based on much more limited information, knowing only the range of price fluctuation, is first proposed by El-Yaniv, Fiat, Karp, and Turpin (2001), as an optimal online trading algorithm is provided by competitive ratio (CR) analysis. They managed to transform the one-way trading problem into a zero-sum game between the trader and a hypothetical adversary, who always tries to inflict worst-case prices on the trader. The value function of the

* Corresponding author. Tel.: +86755 2603 3183/+861062759613.
E-mail addresses: wangw@nankai.edu.cn (W. Wang), wangliying@cueb.edu.cn (L. Wang), ylan@pku.edu.cn, lanyjie@gmail.com (Y. Lan), jxzhang@vcu.edu (J.X. Zhang).

http://dx.doi.org/10.1016/j.ejor.2016.01.031
0377-2217/

game is the ratio between the performance of the online and the offline trading algorithm, and the worst-case ratio for the trader is known as the CR. Because of the complexities evolved from the mathematical form of the value function, the game can hardly be solved directly via backward induction. Instead, El-Yaniv et al. (2001) resorted to an ad-hoc solution procedure based on intuition. After first assuming certain structural properties of the online algorithm, they found the CR and the corresponding online trading algorithm called the threat-based policy, which they proved to be optimal. The research of El-Yaniv et al. (2001) is not without limitations. First, it heavily depends on intuition on the structure of the online algorithm. For more complex problems, such intuition might not be available. More importantly, their intuitive approach has probably only found some very special worst-case scenarios, while missing many other interesting ones. Second, their original threat-based algorithm is unable to improve on the CR whenever the price deviates from the worst-case scenario. Even though they have designed an adaptive online algorithm to fix this issue to some extent, they still depend on the same intuition.

We are the first to apply the mini–max regret criterion proposed by Savage (1951) to solve the one-way trading problem with limited information. The regret of a policy under a specific state refers to the loss due to choosing this policy instead of the optimal one under this particular state. The worst regret of a policy is the maximum regret of this policy under all possible states. The mini–max regret criterion aims at finding the policy with the minimum worst regret, which can be utilized for robust decision-making. This approach based on Savage's regret criterion to analyze online algorithms is called competitive difference (CD) analysis, for its similarity to the CR analysis. In fact, the CD analysis finds an online trading algorithm with the smallest worst-case performance gap from the optimal offline algorithm. In other words, the online trading algorithm found by CD analysis is the closest to the optimal offline algorithm if the distance is measured by the worst-case performance gap. As CD analysis has a simpler mathematical form in the value function than CR analysis, it can obtain closed-form solutions directly via backward induction, without having to guess about the structure of the optimal online policy first. The CD analysis of the one-way trading problem reveals for the first time all possible worst-case scenarios, while providing both the non-adaptive and the adaptive online trading policy at the same time. In addition, numerical experiments show that for revenue-maximization problems, the online trading policy obtained from CD analysis outperforms that from CR analysis, earning higher average revenue with lower standard deviation when the trading horizon is not too short.

This paper is organized as follows. In Section 2, we provide a literature review on topics related to the one-way trading problem, including the time series search problem, the one-way trading problem, and some extended work from the research of El-Yaniv et al. (2001). In Section 3, we introduce the general approach of CD analysis. In Section 4, we find the optimal online trading algorithm for the general multi-period one-way trading problem via CD analysis, and study some nice properties of this online trading algorithm, including its robustness, optimality, and the worst-case scenarios. In Section 5, we conduct numerical performance comparison of our online trading algorithm based on CD analysis with that of El-Yaniv et al. (2001) based on CR analysis, as well as the optimal offline algorithm and the algorithm with distribution information. Finally, we summarize our results and point out some promising directions for future research.

2. Related literature

We review the literature on some problems closely related to the one-way trading problem – in fact, some papers even treat some of them together. On the other hand, we will focus on robust optimization with limited information, and leave out the vast amount of literature on Bayesian approaches, for which we refer the reader to Lippman and McCall (1981).

The time series search problem is closely related to the one-way trading problem, as the only difference between them is that the former only allows for one transaction, while the latter allows for multiple transactions to complete the total trading quantity. Also, the randomized online search can be regarded as a deterministic one-way trading problem and vice versa (El-Yaniv et al., 2001). For the search problem with constant upper and lower price bounds, El-Yaniv et al. (2001) find both the optimal deterministic and randomized reservation price policy. For the search problem with time-varying price bounds, Damaschke, Ha, and Tsigas (2009) design an optimal deterministic algorithm and a nearly-optimal randomized algorithm for the continuous time model, as well as an optimal randomized algorithm for the discrete time model. Xu, Zhang, and Zheng (2011) extend the search problem by introducing a general profit function which increases in price but decreases in time, and they propose two optimal deterministic algorithms respectively for both cases when the search duration is either known or unknown beforehand.

The next related problem is the k-search problem. In a k-search problem, the trader must trade k indivisible items one by one within a horizon, so he should try to find the k highest (or lowest) prices in a sequence. Lorenz, Panagiotou, and Steger (2009) first consider the k-search problem. They figure out the optimal deterministic and randomized algorithms for both the maximization and the minimization objectives. Zhang, Xu, Zheng, and Liu (2011) generalize this research by allowing to trade multiple items at each period and present an optimal deterministic algorithm for the case where the quantity of the item is smaller than the length of the trading horizon. Zhang, Xu, Zheng, and Dong (2012a) further consider the so-called multiple time series search problem in which one product arrives at each period and the storage capacity is limited, and they compare three types of online algorithms.

The one-way trading problem with limited information is first considered by El-Yaniv et al. (2001), where only constant upper and lower price bounds are known. They design a threat-based online trading policy via CR analysis. Chen, Kao, Lyuu, and Wong (2001) as well as Zhang, Xu, Zheng, and Dong (2012b) examine the one-way trading problem with time-varying price bounds which depend on the price at the previous period, and present the optimal deterministic and randomized policy. Larsen and Wohlk (2010) apply the competitive analysis techniques to the online inventory control problem with inventory holding costs and possible order costs. Dai, Jiang, and Feng (2016) improve the upper bound of the CR given by Larsen and Wohlk (2010) in the case with order costs and bounded storage capacity.

Although CR analysis is frequently applied to search and one-way trading problems in the literature, it is often criticized for being too conservative. Several researches try to amend it by adopting alternative performance measures of the online algorithms, which are summarized by Dorrigiv and Lopez-Ortiz (2005). Boyar, Larsen, and Maiti (2014),, compare multiple performance measures for the time series search problem. Fujiwara, Iwama, and Sekiguchi (2011) conduct average-case competitive analysis and derive optimal online algorithms for the one-way currency exchange problem in which the upper bound of prices follows a probability distribution. Mohr and Schmidt (2013) compare the performance of online algorithms based on competitive analyses with different types of information about future prices, and try to measure the value of information. They point out that more information does not necessarily lead to better performance of online algorithms.

All those amendments mentioned above adhere to the CR analysis, and the inherent analytical difficulties prevent them from