A Game Theoretic Approach to Newsvendor Problems with Censored Markovian Demand

Parisa Mansourifard
Ming Hsieh Department of Electrical Engineering
University of Southern California
Los Angeles, California, 90089 USA
parisama@usc.edu

Farzaneh Mansourifard
Department of Engineering & Technology
Alzahra University
Tehran, Iran
farzaneh.mansourifard@gmail.com

Bhaskar Krishnamachari
Ming Hsieh Department of Electrical Engineering
University of Southern California
Los Angeles, California, 90089 USA
bkrishna@usc.edu

Abstract

This paper studies the Newsvendor problem for a setting in which (i) the demand is temporally correlated, (ii) the demand is censored, (iii) the distribution of the demand is unknown. The correlation is modeled as a Markovian process. The censoring means that if the demand is larger than the action (selected inventory), only a lower bound on the demand can be revealed. The uncertainty set on the demand distribution is given by only the upper and lower bound on the amount of the change from a time to the next time. We propose a robust approach to minimize the worst-case total cost and model it as a min-max zero-sum repeated game. We prove that the worst-case distribution of the adversary at each time is a two-point distribution with non-zero probabilities at the extrema of the uncertainty set of the demand. And the optimal action of the decision-maker can have any of the following structures: (i) a randomized solution with a two-point distribution at the extrema, (ii) a deterministic solution at a convex combination of the extrema. Both above solutions balance the over-utilization and under-utilization costs. Finally, we extend our results to uni-model cost functions.

Keywords

Distribution-free newsvendor problem, Markovian process, Min-max, Game theory, Uncertainty set, Robust optimization

1. Introduction

Newsvendor problem or perishable inventory control problem has been a research topic for many years (Arrow 1951). The newsvendor model relates by analogy to the situation faced by a newspaper vendor who must decide on how many newspapers to stock since he doesn't know how many demand (customer) he might have, and he knows that the leftover newspapers cannot be sold the next day (it is perishable in some sense). Since then, different solutions under different assumptions have been presented. One of the approaches to tackle such a challenge is formulating the problem as a robust optimization problem. For a complete literature review on robust optimization and its application in inventory control problems, we refer to Gabrel 2014 and Xin 2015a. Most of the works in the literature focus on the fully observable demand with ambiguous distribution (e.g. See 2010, Xin 2013). Among them, some of the works assume the demand is independent and identically distributed (i.i.d.) at different time periods (e.g. Ding 2002, See 2009, Solyali 2016). Some other works (such as Negoescu 2008, Besbes 2013) consider the case where the demand distribution is i.i.d but unknown, and to solve such a problem, they propose a learning process to estimate the distribution and make the decision. In recent years, it has been observed that the demand distribution is not necessarily i.i.d and it can have correlation over time (Xin 2015b, Carrizosa 2016, Natarajan 2017, Tai 2016). For example, In Xin 2015b, the Martingale demand is considered and the minimax optimal policy is explicitly computed in a closed form. Hu 2016 studied the inventory control problem with Markov-modulated demand. In Carrizosa 2016, a robust approach is proposed for the Newsvendor problem with auto regressive (AR) demand with an unknown distribution. Using numerical experiments, they show that the proposal usually outperforms the previous benchmarks in terms of robustness and the average revenue. The distributionally robust version of the inventory problems over the set of distributions satisfying the known information, which is usually mean and covariance of demand, is studied in Natarajan 2017. The authors show that a three-point distribution achieves the worst-case expected profit and derive a closed-form expression for the problem.

Note that in most of these studies, the demand is assumed to be fully observed. However, there are some research papers which study the inventory control problem with censored (partially observed) and temporally correlated demands (non i.i.d.) (e.g. Lu 2008, and Bisi 2011). In these works, a Bayesian scheme is employed to dynamically update the demand distribution for the newsvendor problem with a storable or perishable inventory. As another example, in Bensoussan 2007, a perishable inventory management problem with a memory (Markovian with known transition probabilities) and partially observable demand process is considered. In our previous work (Mansourifard 2017), we studied a Newsvendor problem with Markovian and censored demand, with the assumption that the transition probabilities are known, as well. In this paper, we extend the work to the case where the transition probability matrix is unknown and only the upper and lower bounds are given.

The contribution of this paper is as follows:

- To our knowledge, this paper is the first work tackling the robust newsvendor problem with temporally correlated demand with censored demand, and we use a game theoretic approach in our solution. We model this problem as a zero-sum repeated game with incomplete information (Sorin 2002, Zamir 1992) and derive the solution in a closed-form.
- We prove that the worst-case distribution of the adversary at each time is a two-point distribution with non-zero probabilities at the lower and upper bound of the uncertainty set.
- The optimal action to minimize the worst-case cost-to-go can have be any of the following two formats: (i) a randomized solution with a two-point distribution at the lower and upper bound of the uncertainty set, (ii) a deterministic solution at a convex combination of the lower and upper bounds of the uncertainty set
- Both the possible solutions balance the over-utilization and under-utilization costs. In other words, if the over-utilization cost is larger than the under-utilization cost, the decision-maker assigns a higher probability to the lower bound (for the solution (i)) or chooses a lower action (for the solution (ii)) to behave conservatively. Otherwise, he behaves more aggressively to increase the chance of getting full observation which can be useful in decreasing the future cost.
- We also show that similar results hold for a more general class of cost functions that are uni-model on the difference between the demand and the action.

2. Problem Formulation

We consider a single-item multi-period Newsvendor problem. The newsvendor model is a mathematical model in operations management and applied economics that is used to decide about the optimal inventory level (action) and it is typically assumes that the prices are fixed and the demand is uncertain for a perishable product. The decision-maker must select the action (e.g. inventory) r_t to satisfy the demand a_t where $t = 1, \dots, T$ is the time step with the finite horizon T. The goal of the decision maker is to minimize the total expected cost over the horizon.

In this paper, we assume the demand a_t is temporally correlated over time as a Markovian random process given by $a_t = a_{t-1} + \delta_t$ with $\delta_t$as a linear transition of the demand from time step $t-1$ to t. In general, we have no information about δ_t, however, we assume that is bounded as $\delta_t \in \{-\delta_t^l, \dots, 0, \dots, \delta_t^h\}$ where $-\delta_t^l$, δ_t^h are the lower and upper bound on the transition, respectively.

As mentioned before, the demand is not necessarily fully observed at each time step. Thus, we consider the case of censored inventory problem in which at each time t, if $r_t > a_t$, the decision-maker gets a full observation about a_t, and if $r_t \leq a_t$, only partial observation about a_t reveals (i.e. a_t is censored).

For a given demand a_t and a selected action r_t, the decision maker faces an immediate cost as:

$$C(a_t, r_t) = \begin{cases} c_u(r_t - a_t) & if\ r_t > a_t \\ c_l(a_t - r_t) & if\ r_t \leq a_t \end{cases} \tag{1}$$

where c_u and c_l are over-utilization and under-utilization cost coefficients, respectively. The goal is to minimize the total expected cost accumulated over the finite horizon. Since the demand is unknown, this goal could be formulated as a min-max optimization problem:

$$C_1^*\,(r_1^l\,, r_1^h) = \min_{P_{r_1},\dots,P_{r_T}} \max_{\mathbb{P}_{a_1},\dots,\mathbb{P}_{a_1}} \sum_{t=1}^{T} \mathbb{E}_{r_t}\mathbb{E}_{a_t}[C(a_t, r_t)|F_t], \tag{2}$$

where F_t is the information available to the player before time t, $\mathbb{p}_{r_t}$ and $\mathbb{p}_{a_t}$ are probability distribution functions (PDFs) of the action and the demand at time t, respectively. Since the transition of the demand is bounded, the action would also be bounded in $[r_t^l\ , r_t^h]$. Let $C_t^*\ (r_t^l\ , r_t^h)$ indicate the min-max expected cost-to-go from time t onward where $C_1^*\ (r_1^l\ , r_1^h)$ is given by (2).

After taking the action r_t and the observation revealed about the demand, the bounds on the actions can be updated for the next time step:

$$r_{t+1}^l = \begin{cases} a_t - \delta_t^l & if\ r_t > a_t \\ r_t - \delta_t^l & if\ r_t \leq a_t \end{cases}$$

$$r_{t+1}^h = \begin{cases} a_t + \delta_t^h & \text{if } r_t > a_t \\ r_t^h + \delta_t^h & \text{if } r_t \leq a_t \end{cases}$$

Now, the goal is to find the best actions r_t^* that achieves the min-max at (2). Fig. 1 shows an example of the demand path and the sequence of the taken actions with the corresponding costs and the bounds on the actions.

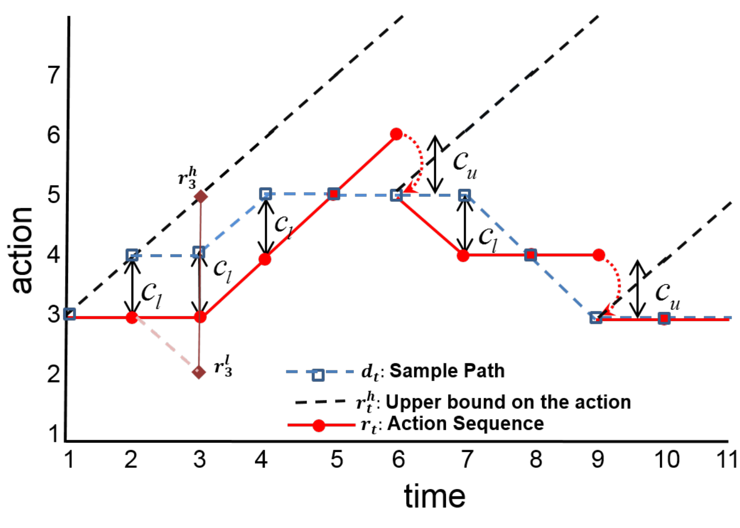

Fig. 1. An example of the demand path and the sequence of the taken actions.

3. The Game Theoretic Approach

This can be modeled as a game between the adversary and the decision-maker. The sufficient statistic for F_t at each time step t, is the support $\{r_t^l, r_t^h\}$ and the adversary chooses the probability distribution of $a_t \in \{r_t^l, r_t^h\}$ to maximize the expected cost-to-go for the selected distribution of the action r_t. The solutions are given in the following theorem (and we prove them using induction):

Theorem 1-a) The worst-case distributions $\mathbb{p}_{at}$ are two-point distributions with non-zero probabilities at r_t^l and r_t^h, for all $t = 1, \dots, T$. And there are two possible solution to the min-max problem:

1. $r_t^* = \frac{c_u r_t^l + y_t r_t^h}{c_u + y_t}$, $p_{r_t}(r_t^*) = 1$, and p_{a_t} can be any two-point distribution at extrema
2. $p_{r_t}(r_t^l) = 1 - p_{r_t}(r_t^h) = \frac{c_u}{c_u + c_l}$, and

$$p_{a_t}(a_t^l) = 1 - p_{a_t}(a_t^h) = \frac{y_t}{c_u + y_t},$$

where,

$$y_t = c_l + \frac{c_u y_{t+1}}{c_u + y_{t+1}}, \forall t = 1, \dots, T-1$$

$$y_T = c_l.$$

Theorem 1-b) The min-max cost-to-go at time t is obtained as:

$$C_t^*(r_t^l, r_t^h) = C_t^*(r_t^h - r_t^l) = \Delta_{t+1} + \frac{c_u y_t}{c_u + y_t}(r_t^h - r_t^l),$$

where,

$$\Delta_t = \Delta_{t+1} + \frac{c_u y_t}{c_u + y_t}(\delta_t^h + \delta_t^l), \forall t = 1, \dots, T-1,$$

$$\Delta_t = \frac{c_u y_t}{c_u + y_t}(\delta_t^h + \delta_t^l).$$

Proof:

We use induction to prove the both parts of the Theorem 1. First, at horizon T, we need to solve the single-period version of this problem:

$$\min_{\mathbb{P}_{r_T}} \max_{\mathbb{P}_{a_T}} \mathbb{E}_{r_T} \mathbb{E}_{a_T}[C(a_T, r_T)] = \mathbb{E}_{r_T}[\int_{x=r_T^l}^{r_T} c_u(r_T - x) p_{a_T}(x) dx + \int_{x=r_T}^{r_T^h} c_l(x - r_T) p_{a_T}(x) dx].$$

The PDF maximizing the above equation is a two-point distribution where only $p_{a_T}(r_T^l)$ and $p_{a_T}(r_T^h)$ are non-zero. Fig. 2 shows the graph of the cost functions for all a and any two pairs of $r, r' \in [r_T^l, r_T^h]$. As it is obvious, for any pair of r, r', the values of a that can affect the min-max cost are $a = r_T^l$ and $a = r_T^h$ since their corresponding cost functions are included in the upper bound of the costs. Therefore, we can ignore any adversary actions inside the bounds, i.e. $a \in (r_T^l, r_T^h)$ which confirms that the distribution is two-point at horizon.

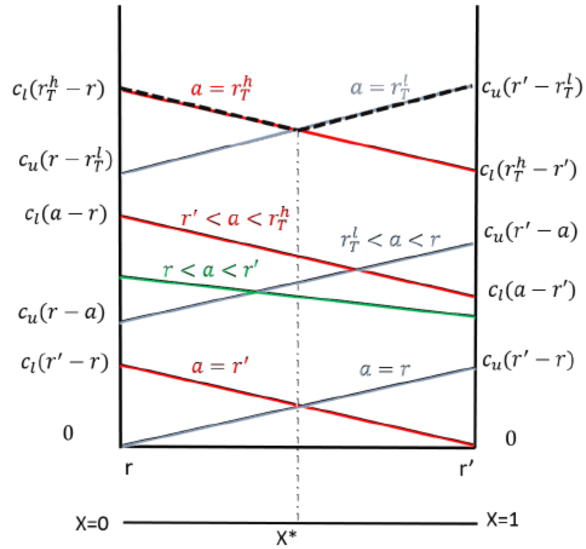

Fig.2. The cost functions $XC_T(a,r) + (1-X)C_T(a,r')$; WLOG assume $r < r', X = \frac{p_{r_T}(r')}{p_{r_T}(r)+p_{r_T}(r')}$

Now to find the best action of the decision-maker at horizon, we plot the graph of the cost functions for different values of r in Fig. 3, and find the randomized action with the best probability density of $p_{a_T}(r_T^l) = x$ and $p_{a_T}(r_T^h) = 1 - x$ or the best deterministic action r_T^*.

For both possible actions, the expected cost can be computed as:

$$C_T^*(r_T^l, r_T^h) = \frac{c_l c_u}{c_l + c_u}(r_T^h - r_T^l).$$

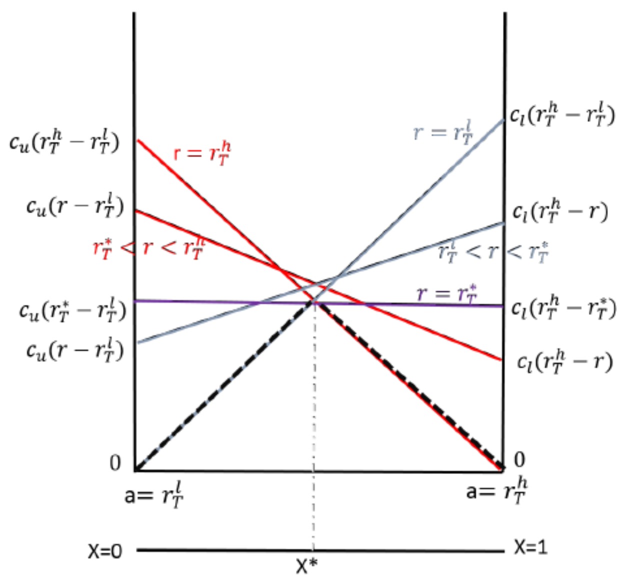

Fig. 3. The cost functions $XC_T(r_T^l,r)+(1-X)C_T(r_T^h,r)$; $X = \frac{p_{a_T}(r_T^h)}{p_{a_T}(r_T^h)+p_{a_T}(r_T^l)}$. The optimal distribution of player could have non-zero probabilities at $r \in \{r_T^l, r_T^*, r_T^h\}$

Now for time steps $t = 1, \ldots, T-1$, we use induction to find the best action distributions (or best deterministic action) and the worst-case adversary distributions. If the solutions are valid for $t + 1$, for time t we have:

$$\min_{\mathbb{P}_{r_T}} \max_{\mathbb{P}_{a_T}} \mathbb{E}_{r_t}[\int_{x=r_t^l}^{r_t} c_u(r_t - x)p_{a_t}(x)dx + \int_{x=r_t}^{r_t^h} c_l(x - r_t)p_{a_t}(x)dx$$

$$+ \int_{x=r_t^l}^{r_t} p_{a_t}(x)C_{t+1}^*(x - \delta_{t+1}^l, x + \delta_{t+1}^h)dx$$

$$+ \int_{x=r_t}^{r_t^h} p_{a_t}(x)dx. C_{t+1}^*(r_t - \delta_{t+1}^l, r_t^h + \delta_{t+1}^h)]$$

Since $C_{t+1}^*(y,z) = C_{t+1}^*(z-y) = \Delta_{t+2} + \frac{c_u y_{t+1}}{c_u + y_{t+1}}(z-y)$,

$$\min_{\mathbb{P}_{r_t}} \max_{\mathbb{P}_{a_t}} \mathbb{E}_{r_t}[\int_{x=r_t^l}^{r_t} c_u(r_t - x)p_{a_t}(x)dx + \int_{x=r_t}^{r_t^h} c_l(x - r_t)p_{a_t}(x)dx$$

$$+ \int_{x=r_t^l}^{r_t} p_{a_t}(x)dx. (\Delta_{t+2} + \frac{c_u y_{t+1}}{c_u + y_{t+1}}(\delta_{t+1}^l + \delta_{t+1}^h))$$

$$+ \int_{x=r_t}^{r_t^h} p_{a_t}(x)dx. (\Delta_{t+2} + \frac{c_u y_{t+1}}{c_u + y_{t+1}}(r_t^h - r_t + \delta_{t+1}^l + \delta_{t+1}^h))]$$

$$= \Delta_{t+1} + \min_{\mathbb{P}_{r_t}} \max_{\mathbb{P}_{a_t}} \mathbb{E}_{r_t}[\int_{x=r_t^l}^{r_t} c_u(r_t - x)p_{a_t}(x)dx + \int_{x=r_t}^{r_t^h} c_l(x - r_t)p_{a_t}(x)dx$$

$$+ \int_{x=r_t}^{r_t^h} p_{a_t}(x)dx. \frac{c_u y_{t+1}}{c_u + y_{t+1}}(r_t^h - r_t)]$$

In other words,

$$C_t^*(r_t^l, r_t^h) = \Delta_{t+1} + \min_{\mathbb{P}_{r_t}} \max_{\mathbb{P}_{a_t}} \mathbb{E}_{r_t}\mathbb{E}_{a_t}[C_t'(a_t, r_t)],$$

where,

$$C_t'(a_t, r_t) = \begin{cases} c_u(r_t - a_t) & if\ r_t > a_t \\ c_l(a_t - r_t) + \frac{c_u y_{t+1}}{c_u + y_{t+1}}(r_t^h - r_t) & if\ r_t \le a_t \end{cases}$$

As it is shown in Fig. 4, we can ignore the adversary actions of $a \in (r_t^l, r_t^h)$. Now to find the best distributions of the decision-maker, in Fig. 5 we plot the graph of the cost functions for different values of r and find the best probability density for $p_{a_t}(r_t^h) = X$ and $p_{a_t}(r_t^l) = 1 - X$.

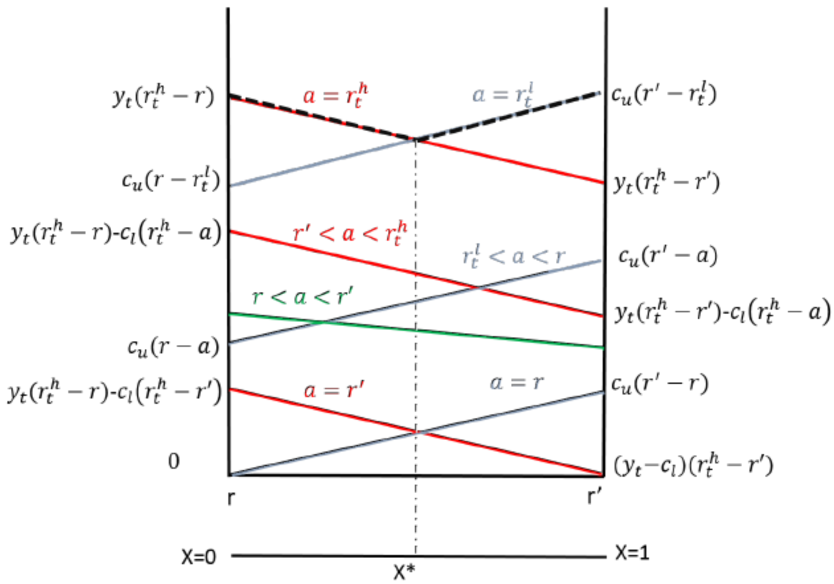

Fig. 4. The cost functions $XC_t'(a,r) + (1-X)C_t'(a,r')$; WLOG assume $r < r'$, $X = \frac{p_{r_t}(r')}{p_{r_t}(r)+p_{r_t}(r')}$ Note that to get the actual cost function we should add all of them by Δ_{t+1}.

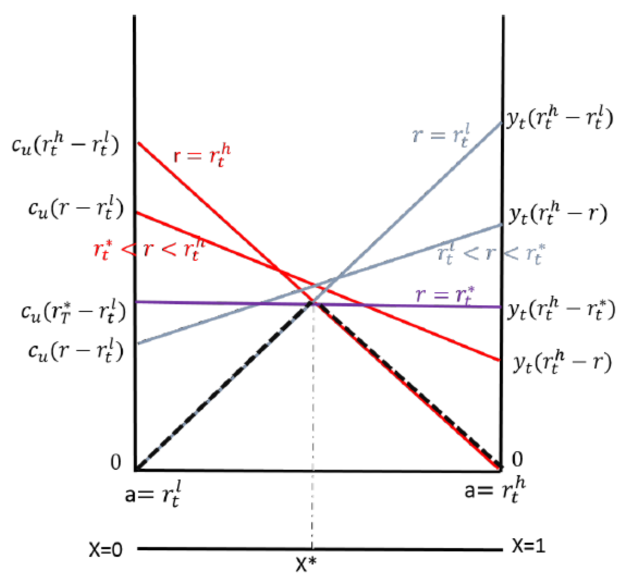

Fig. 5. The cost functions $XC_t'(r_t^l,r) + (1-X)C_t'(r_t^h,r)$; $X = p_{a_t}(r_T^h) = 1 - p_{a_t}(r_t^l)$. The optimal distribution of player could have non-zero probabilities at $r \in \{r_t^l, r_T^*, r_t^h\}$.

And the worst-case distribution is a two-point distribution at r_t^l and r_t^h. Therefore:

$$C_t^*(r_t^l, r_t^h) = p_{a_t}(r_t^l)c_u(r_t^* - r_t^l) + (1 - p_{a_t}(r_t^l))c_l(r_t^h - r_t)$$

$$+p_{a_t}(r_t^l)\left[\Delta_{t+2} + \frac{c_u y_{t+1}}{c_u + y_{t+1}}(\delta_{t+1}^l + \delta_{t+1}^h)\right]$$

$$+\left(1-p_{a_t}(r_t^l)\right)\left[\Delta_{t+2}+\frac{c_u y_{t+1}}{c_u+y_{t+1}}(r_t^h-r_t^*+\delta_{t+1}^l+\delta_{t+1}^h)\right]$$

$$=p_{a_t}(r_t^l)c_u(r_t^*-r_t^l)+\left(1-p_{a_t}(r_t^l)\right)\left(c_l+\frac{c_u y_{t+1}}{c_u+y_{t+1}}\right)(r_t^h-r_t)+\Delta_{t+2}$$

$$=p_{a_t}(r_t^l)c_u(r_t^*-r_t^l)+\left(1-p_{a_t}(r_t^l)\right)y_t(r_t^h-r_t)+\Delta_{t+2}$$

$$=c_u(r_t^*-r_t^l)+\Delta_{t+2}.$$

which results in the optimal actions given in Theorem 1-a and the min-max expected cost given in Theorem 1-b, thus completes the proof of Theorem 1.

4. Extension of Cost Function

We can get similar results for general form of uni-modal cost functions, given by:

$$C(a_t,r_t)=\begin{cases}C_u(r_t-a_t) & if\ r_t>a_t\\ C_l(a_t-r_t) & if\ r_t\le a_t\end{cases}$$

where $C_u(y)$and $C_l(y)$ are increasing functions of y.

Lemma 1-a) The worst-case distribution at all time steps $t = 1,\dots,T$ are two-point distributions $p_{a_t}(r_t^l)\neq 0$ and $p_{a_t}(r_t^h)\neq 0$.

Lemma 1-b) The min-max expected cost has the following property: $C_t^*(y+x,z+x) = C_t^*(y,z)$.

Proof:

The min-max cost-to-go at time t is given by:

$$C_t^*(r_t^l,r_t^h)=\min_{r_t}\max_{\mathbb{P}_{a_t}}\int_{x=r_t^l}^{r_t}p_{a_t}(x)[c_u(r_t-x)+C_{t+1}^*(x-\delta_t^l,x+\delta_t^h)]\,dx$$

$$+\int_{x=r_t}^{r_t^h}p_{a_t}(x)[c_l(x-r_t)+C_{t+1}^*(r_t-\delta_t^l,r_t^h+\delta_t^h)]\,dx$$

At horizon T the worst-case distribution is a two-point distribution and:

$$r_T^*=\{r: C_u(r-r_T^l)=C_l(r_T^h-r)\}$$

and the expected cost equals:

$$C_T^*(r_T^l,r_T^h)=C_l(r_T^h,r_T^*)=C_u(r_T^*,r_T^l).$$

This shows that Lemma 1 is true at $t = T$, now if it is true at $t+1$, for time t we have:

$$C_t^*(r_t^l,r_t^h)=\min_{r_t}\max_{\mathbb{P}_{a_t}}\int_{x=r_t^l}^{r_t}p_{a_t}(x)[c_u(r_t-x)+C_{t+1}^*(r_t^l-\delta_t^l,r_t^l+\delta_t^h)]\,dx$$

$$+\int_{x=r_t}^{r_t^h}p_{a_t}(x)[c_l(x-r_t)+C_{t+1}^*(r_t-\delta_t^l,r_t^h+\delta_t^h)]\,dx$$

where we replace x at $C_{t+1}^*(x-\delta_t^l,x+\delta_t^h)$ with r_t^l. From the above equation, it is obvious that: (i) the worst-case distribution is a two-point distribution, and (ii) if we add a fixed value to r_t^l and r_t^h the minimizing r_t will be added with the same amount and thus C_t^* (r_t^l+x; r_t^h +x) = C_t^* (r_t^l, r_t^h) for any x.

And recursively:

$$r_t^* = \{r: C_l(r_t^h - r) + C_{t+1}^*(r - \delta_t^l, r_t^h + \delta_t^h) = C_u(r - r_t^l) + C_{t+1}^*(r_t^l - \delta_t^l, r_t^l + \delta_t^h)\}$$

or a randomized solution as follows:

$$p_{r_t}(r_t^l) = 1 - p_{r_t}(r_t^h) = \frac{r_t^h - r_t^*}{r_t^h - r_t^l},$$

Or any combination of non-zero probabilities at $r_t \in \{r_t^l, r_t^*, r_t^h\}$ which proves Lemma 1-a. And the expected cost equals:

$$C_t^*(r_t^l, r_t^h) = C_l(r_t^h - r_t^*) + C_{t+1}^*(r_t^* - \delta_t^l, r_t^h + \delta_t^h) = C_u(r_t^* - r_t^l) + C_{t+1}^*(r_t^l - \delta_t^l, r_t^l + \delta_t^h).$$

This proofs Lemma 1-b.

5. Conclusion

We have studied the Newsvendor problem with the following challenges: (i) the demand is temporally correlated as a Markovian process, (ii) the demand can only be censored (i.e. partially observable), (iii) the distribution of the demand and the transition probabilities of the Markovian process are unknown and only upper and lower bounds on the transitions are given. We modeled this problem as a min-max zero-sum repeated game. We have proved that the worst-case distribution of the adversary at each time is a two-point distribution with non-zero probabilities at the lower and upper bound of the uncertainty set. The optimal action to minimize the worst-case cost-to-go can have be any of the following two formats: (i) a randomized solution with a two-point distribution at the lower and upper bound of the uncertainty set. If the over-utilization cost is larger than the under-utilization cost, higher probability is assigned to the lower bound to behave conservatively. Otherwise, higher probability is assigned to the upper bound to behave more aggressively and increase the chance of full observation. (ii) a deterministic solution at a convex combination of the lower and upper bounds of the uncertainty set, which also balance the over-utilization and under-utilization costs. Finally, we showed that similar results hold for a more general class of cost functions that are uni-model on the difference between the demand and the action.

Acknowledgements

This work was supported in part by the U.S. National Science foundation under ECCS-EARS awards numbered 1247995 and 1248017, by the Okawa foundation through an award to support research on “Network Protocols that Learn”. Parisa Mansourifard was supported by AAUW American Dissertation Completion Fellowship for 2015-2016.

References

Arrow, K. J., Harris, T., and Marschak, J., “Optimal inventory policy,” Econo-metrica: Journal of the Econometric Society, pp. 250–272, 1951.

Zamir, S., “Repeated games of incomplete information: Zero-sum”, Handbook of Game Theory with Economic Applications, vol. 1, pp. 109-154, 1992.

Ding, X., Puterman, M. L., Bisi, A., “The censored newsvendor and the optimal acquisition of information”, Operations Research, vol. 50, no. 3, pp. 517–527, 2002.

Sorin, S., “A first course on zero-sum repeated games”, Springer Science & Business Media, vol. 37, 2002.

Negoescu, D., Frazier, P., Powell, W., “Optimal learning policies for the newsvendor problem with censored demand and unobservable lost sales”, URL: http://people.orie.cornell.edu/pfrazier/pub/learning-newsvendor.pdf, 2008.

Lu, X., Song, J., Zhu, K., “Analysis of perishable-inventory systems with censored demand data”, Operations Research, vol. 56 no. 4, pp. 1034-1038, 2008.

Bensoussan, A., Akanyıldırım, M. C¸ Sethi, S. P., “A multiperiod newsvendor problem with partially observed demand”, Mathematics of Operations Research, vol. 32, no. 2, pp. 322–344, 2007.

Bisi, A., Dada, M., Tokdar, S., “A censored-data multi-period inventory problem with newsvendor demand distributions”, Manufacturing & Service Operations Management, vol. 13, no. 4, pp. 525-533, 2011.

See, C.-T. and Sim, M., “Robust approximation to multiperiod inventory management”, Operations research, vol. 58, no. 3, pp. 583-594, 2010.

Besbes, O., Muharremoglu, A., “On implications of demand censoring in the newsvendor problem”, Management Science, vol. 59, no. 6, pp. 1407-1424, 2013.

Gabrel, V., and Murat, C., and Thiele, A., "Recent advances in robust optimization: An overview", European journal of operational research, vol. 235, no. 3, pp. 471-483, 2014.

Xin, L., "New approaches to inventory control: algorithms, asymptotics and robustness", PhD. Thesis Georgia Institute of Technology, 2015a.

Xin, L., and Goldberg, D. A., "Distributionally robust inventory control when demand is a martingale", arXiv preprint arXiv:1511.09437, 2015b.

Solyali, O., and Cordeau, J.-F., and Laporte, G., "The impact of modeling on robust inventory management under demand uncertainty", Management Science, vol. 62, no. 4, pp. 1188-1201, 2015.

Park, K., "Min-Max Stochastic Optimization with Applications to the Single-Period Inventory Control Problem", Management Science & Financial Engineering, vol. 21, no. 1, 2015.

Carrizosa, E., and Olivares-Nadal, A. V. and Ramrez-Cobo, P., "Robust newsvendor problem with autoregressive demand", Computers & Operations Research, vol. 68, pp. 123-133, 2016.

Mansourifard, P., Javidi, T., Krishnamachari, B., "Percentile Policies for Tracking of Markovian Random Processes with Asymmetric Cost and Observation", arXiv preprint arXiv:1703.01261, 2017.

Tai, A. H, Ching, WK, "Recent advances on Markovian models for inventory research", International Journal of Inventory Research, vol. 3, no. 3, pp. 198-216, 2016.

Hu, J., Zhang, Ch., Zhu, Ch., "(s, S) Inventory Systems with Correlated Demands", INFORMS Journal on Computing, vol. 28, no. 4, pp. 603-611, 2016.

Natarajan, K., and Sim, M., and Uichanco, J., "Asymmetry and ambiguity in newsvendor models", Management Science, 2017.

Qiu, R., and Sun, M., and Lim, Y. F., "Optimizing (s, S) policies for multi-period inventory models with demand distribution uncertainty: Robust dynamic programing approaches", European Journal of Operational Research, vol. 261, no. 3, pp. 880-892, 2017.

Biographies

Parisa Mansourifard received the B.S. and M.S. in electrical engineering from Sharif university of technology, Tehran, Iran, in 2008 and 2010 respectively. She also the M.S. in computer science and Ph.D. in electrical engineering from University of Southern California, Los Angeles, CA, USA, in 2015 and 2017, respectively. During her Ph.D. she held Viterbi Dean fellowships in 2011-2014 and AAUW dissertation completion fellowship in 2015-2016. She is currently a data scientist at Supplyframe Inc. and a part-time lecturer at University of Southern California. Her research interest is decision-making under uncertainty, machine learning and stochastic optimization.

Farzaneh Mansourifard is a master student at department of engineering and technology at Alzahra university, Tehran, Iran. She earned B.S. in industrial engineering from Tabriz university, Tabriz, Iran, 2016. Her research interest is optimization with the applications in healthcare and supply chain.

Bhaskar Krishnamachari (M02 – SM14) received the B.E. degree in electrical engineering at The Cooper Union, New York, NY, USA, in 1998, and the M.S. and Ph.D. degrees from Cornell University, Ithaca, NY, USA, in 1999 and 2002, respectively. He is currently Professor and Ming Hsieh Faculty Fellow in the Department of Electrical Engineering at the Viterbi School of Engineering, University of Southern California, Los Angeles, CA, USA. He is also the Director of the USC Viterbi School of Engineering Center on Cyber-Physical Systems and the Internet of Things. His primary research interest is in the design and analysis of algorithms, protocols, and applications for next-generation wireless networks.