Solving the Newsvendor Problem using Stochastic Approximation: A Kiefer-Wolfowitz Algorithm Approach

Quan Yuan[*], Zhixin Yang, Yayuan Xiao

Department of Mathematical Sciences, Ball State University, Muncie IN, USA
*Corresponding author: qyuan@bsu.edu

Abstract   This paper investigates the application of the Kiefer-Wolfowitz (KW) algorithm, a stochastic approximation technique, to solve the newsvendor problem under uncertain demand. The proposed approach enables the newsvendor to learn from observed profits and converge to the optimal order quantity, even when the demand distribution is unknown. Numerical experiments demonstrate the algorithm's effectiveness in handling stochastic demand and provide insights into its convergence properties. The paper highlights the potential of stochastic approximation methods in tackling inventory management challenges and discusses future research directions.

Keywords: *Newsvendor problem, stochastic approximation, Kiefer-Wolfowitz algorithm*

Cite This Article: Quan Yuan, Zhixin Yang, and Yayuan Xiao, "Solving the Newsvendor Problem using Stochastic Approximation: A Kiefer-Wolfowitz Algorithm Approach." *American Journal of Applied Mathematics and Statistics*, vol. 12, no. 2 (2024): 24-27. doi: 10.12691/ajams-12-2-1.

1. Introduction

1.1. Introduction of Newsvendor Problem

The newsvendor problem, also known as the single-period inventory management problem, is a classic model in operations research and management science. It describes a scenario where a decision-maker needs to determine the optimal order quantity of a product with uncertain demand to maximize expected profit. The problem can be formulated as follows.

A newsvendor needs to order newspapers in advance from the post office and determine the quantity of newspapers to order, denoted as x copies, at a price of c_v per copy. The selling price of each newspaper is known to be p per copy. If the newsvendor does not sell all the newspapers on that day, the recycling center will buy back the newspapers at a price of s_v per copy. Let the daily demand for newspapers be ξ. If $x > \xi$, the daily remainder of newspapers is $x - \xi$; otherwise, it is 0. Thus, the newsvendor's profit is:

$$f(x,\xi) = \begin{cases} (p - c_v)x, & x \le \xi, \\ (s_v - c_v)x + (p - s_v)\xi, & x > \xi. \end{cases} \quad (1)$$

In practical problems, the demand for newspapers ξ is usually a random variable, which leads to the profit function $f(x,\xi)$ also being a random variable. Since it is not possible to accurately predict the actual profit from ordering ξ newspapers, a natural method is to consider the expected profit:

$$E[f(x,\xi)] = \int_0^x [(s_v - c_v)x + (p - s_v)\xi]\phi(\xi)d\xi + \int_x^\infty (p - c_v)x\phi(\xi)d\xi \quad (2)$$

where $\phi(\xi)$ represents the probability density function of the demand ξ. The newsvendor problem is to find the optimal order quantity x that maximizes the expected profit $E[f(x,\xi)]$.

The newsvendor model has been widely applied in various industries, such as retail, fashion, and perishable goods, to guide inventory decisions. Extensive research has been conducted on the newsvendor problem since its introduction. Arrow et al. [1] laid the foundation for the newsvendor model by deriving the optimal order quantity under stochastic demand. Subsequently, numerous extensions and variations of the model have been proposed to capture different real-world settings. For example, Khouja [2] provided a comprehensive review of the single-period inventory problem and its extensions, considering factors such as pricing, discounting, and supplier credits. Qin et al. [3] reviewed the newsvendor problem under supply chain coordination, addressing issues like contracts, information sharing, and risk management.

1.2. Introduction of Stochastic Approximation

The original work in recursive stochastic algorithms was by Robbins and Monro, who developed and analyzed a recursive procedure for finding the root of a real-valued function $h(x)$ of a real variable x. The function is not known, only noise-corrupted observations could be taken at any values of x selected by the experimenter. If the analytic form of $h(\cdot)$ were known and continuously differentiable, then the problem would be classical in numerical analysis, and many procedures, such as Newton's method, could be used. However, in many practical problems, people do not know the form of the nonlinear function in the equation and can only observe the function values with random errors at given values of the independent variable. In this case, how to find the zero point x_0 of this unknown nonlinear function is an important problem from both practical and theoretical perspectives. If x_n is the value of the independent variable taken at the n^{th} measurement, then the observed value of the function is:

$$Y_n = h(x_n) + \epsilon_n \qquad (3)$$

where $\{\epsilon_n\}$ is a sequence of measurement errors that may depend on x_n. In 1951, Robbins and Monro first proposed and studied a stochastic approximation algorithm. They took the sequence $\{a_n\}$ as:

$$a_i > 0, \ \sum_{i=1}^{\infty} a_i = \infty, \ \sum_{i=1}^{\infty} a_i^2 < \infty \qquad (4)$$

The n^{th} approximation to the zero point x_0 is:

$$x_{n+1} = x_n + a_n Y_n \qquad (5)$$

where Y_n is obtained from (3). This is the well-known Robbins-Monro (RM) algorithm.

In the year following the proposal of the RM algorithm, the Kiefer-Wolfowitz (KW) algorithm appeared, which is an algorithm for finding the extreme value of $h(x)$. The basic idea of the KW algorithm is similar as the RM algorithm. In essence, it seeks the zero point of the derivative of $h(x)$. The difference quotient of the measured values of $h(x)$ is used to estimate the value of the derivative of $h(x)$. Therefore, compared with the RM algorithm, except for some changes in the measurement points, the basic form of the algorithm and the basic ideas of theoretical research are consistent.

Stochastic approximation problems have consistently attracted the attention of many applied mathematicians, statisticians, and systems and control experts. The involved fields include process statistics, system identification and parameter estimation, adaptive control, stochastic optimization and decision-making, and signal processing. Relevant papers and monographs can be referred to [6] and its references.

This paper aims to tackle the newsvendor problem using stochastic approximation, a powerful technique for solving stochastic optimization problems. The key advantage of stochastic approximation is its ability to handle uncertain parameters, making it suitable for dynamic inventory control settings.

The remainder of this paper is organized as follows. Section 2 shows the motivation and implementation of applying the stochastic approximation method to the newsvendor problem. Section 3 provides numerical experiments to demonstrate the effectiveness of the proposed approach. Finally, Section 4 concludes the paper and discusses future research directions.

2. Applying Stochastic Approximation to the Newsvendor Problem

As discussed in the Introduction, the common approach to solving the Newsvendor Problem is to find the maximum value of equation (2). When $\phi(\xi)$ is known and simple, we can obtain an analytical solution. When $\phi(\xi)$ is complex, we can use Monte-Carlo integration or genetic algorithms to calculate the integral and solve (2). However, in most real-life situations, $\phi(\xi)$ is unknown. In such cases, we cannot compute the integral in (2) and can only calculate the "contaminated" profit from (1). This is precisely the type of problem that stochastic approximation algorithms excel at solving. Motivated by this, we will employ the Kiefer-Wolfowitz (KW) algorithm to find the optimal solution.

The stochastic approximation algorithm for solving the Newsvendor Problem is as follows:

Step 0: Estimate the number of newspapers to order on the first day, x_1, based on experience.

Repeat:

Step 1: Let x_n be the number of newspapers ordered on the n^{th} day. Calculate:

$$Y_n = -\frac{[f(x_n+1, \xi) - f(x_n-1, \xi)]}{2} \qquad (6)$$

Step 2: Calculate the number of newspapers to order on the $(n+1)^{\text{th}}$ day:

$$x_{n+1} = x_n + a_n Y_n \qquad (7)$$

End Repeat.

In Step 1, Y_n is the difference estimate of the derivative of $f(x, \xi)$:

$$f'(x, \xi) \approx -\frac{[f(x_n+c, \xi) - f(x_n-c, \xi)]}{2c} \qquad (8)$$

We take $c = 1$ here because the number of newspapers ordered must be an integer. Since we know the profit for x_n, it is very easy to calculate the profit for ordering one more and one less newspaper order.

3. Numerical Examples

We consider several scenarios in this section.

Example 1: We consider the newspaper selling price $p = \$1.00$, buy back price $s_v = 0$, and order price $c_v = 0.25$. If the daily demand for newspapers ξ follows a uniform distribution on [0, 40], calculating (2), we can obtain the best order number $x^* = 30$. Using the KW algorithm and setting $a_n \equiv 1/10$, we obtain the graph shown in Figure 1.
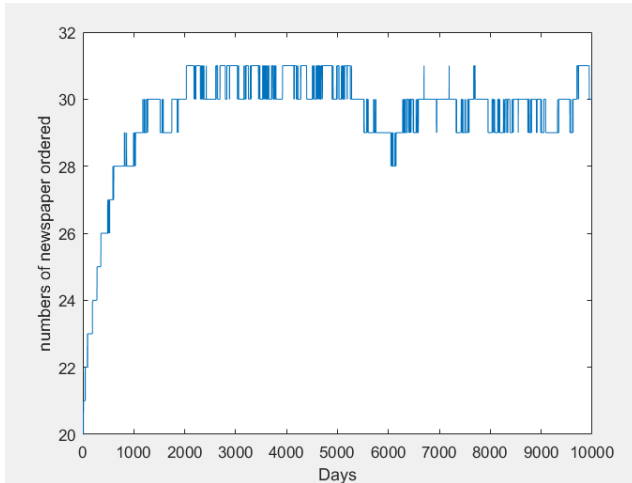
Figure 1. Convergence of the algorithm when demand follows a uniform distribution

As evident from Figure 1, the algorithm gradually converges to the vicinity of $x^* = 30$ and fluctuates between 29 and 31. The average value is 29.95.

Remark: The graph reveals that the convergence speed of x to the optimal value is relatively slow. In practical situations, the following method can be employed to accelerate convergence: after collecting demand quantities ξ for approximately 100 days, we can use the bootstrap method to simulate the entire distribution of ξ through resampling, and then substitute the results into the algorithm for iteration.
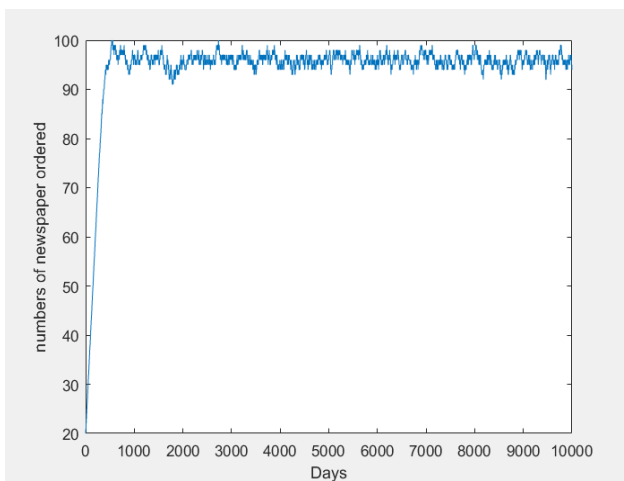
Figure 2. Convergence of the algorithm when demand follows a normal distribution

Example 2: We consider the selling price $p = \$6.00$, buy back price $s_v = \$0.2$, and order price $c_v = \$4.00$. If ξ follows a normal distribution $N(100, 10)$, the optimal

solution given in the literature [7] is $x^* = 96$. Using the KW algorithm and setting $a_n \equiv 1/10$, we obtain the graph shown in Figure 2.

The graph illustrates that the algorithm gradually converges to the vicinity of $x^* = 96$. The average value is 95.87.

4. Conclusion

In this paper, we have demonstrated the effectiveness of applying stochastic approximation algorithms, specifically the Kiefer-Wolfowitz (KW) algorithm, to solve the newsvendor problem under uncertain demand. The proposed approach offers a practical solution for determining optimal order quantities in real-world settings where the demand distribution is unknown or complex.

Through numerical examples, we have shown that the KW algorithm can successfully converge to the optimal order quantity, even when the demand follows various probability distributions. The algorithm's ability to handle stochastic demand and learn from observed profits makes it a valuable tool for dynamic inventory management.

However, our numerical experiments also reveal that the convergence speed of the algorithm can be relatively slow in certain cases. This highlights the need for further research on accelerating the convergence of stochastic approximation algorithms in the context of the newsvendor problem. Potential avenues for improvement include incorporating bootstrap methods to simulate demand distributions and using the results to guide the algorithm's iterations.

Moreover, theoretical analysis of the convergence properties and performance guarantees of the KW algorithm in the newsvendor setting is an important direction for future research. Establishing rigorous convergence results and deriving bounds on the algorithm's performance would provide a stronger foundation for its application in practice.

In conclusion, this paper opens up new possibilities for tackling the newsvendor problem under uncertainty using stochastic approximation techniques. The promising results obtained through numerical experiments motivate further exploration of this approach, both in terms of theoretical analysis and the development of more efficient algorithms. By advancing research in this direction, we can enable better inventory decision-making in a wide range of industries facing stochastic demand.

ACKNOWLEDGEMENTS

References

[1]   Arrow, K.J., Harris, T., Marshak, J., "Optimal Inventory Policy", *Econometrica,* 19(3). 250-272. Jul. 1951.

[2]   Khouja, M., "The single-period (news-vendor) problem: literature review and suggestions for future research", *Omega*, 27(5). 537-

553. Oct. 1999.

[3] Qin, Y., Wang, R., Vakharia, A.J., Chen, Y., Seref, M.M., "The Newsvendor Problem: Review and Directions for Future Research", *European Journal of Operational Research*, 213(2). 361-374. Sep. 2011.

[4] Robbins, H., Monro, S., "A Stochastic Approximation Method", *The Annals of Mathematical Statistics*, 22(3). 400-407. Sep. 1951.

[5] Kiefer, J., Wolfowitz, J., "Stochastic Estimation of the Maximum of s Regression Function", *The Annals of Mathematical Statistics*, 23(3). 462-466. Sep. 1952.

[6] Kushner, H.J., Yin, G., *Stochastic Approximation and Recursive Algorithms and Applications*. Springer-Verlag, 2003.

[7] Gen M., Lin B., "Evolution program for production plan problem". *Engineering Design and Automation*, 1(3). 199-204, Jan. 1995.