

# Warranty Optimization in a Dynamic Environment

**Nedialko B. Dimitrov**  
University of Texas at Austin  
Austin, Texas, USA  
*ned.dimitrov@gmail.com*

**Stefanka Chukova**  
Victoria University of Wellington  
Wellington, NZ  
*stefanka@mcs.vuw.ac.nz*

## Abstract

A product warranty is an agreement offered by a producer to a consumer to replace or repair a faulty item, or to partially or fully reimburse the consumer in the event of a failure. Warranties are very widespread and serve many purposes, including protection for producer, seller, and consumer. They are used as signals of quality and as elements of marketing strategies. In this study we review the notion of an online convex optimization algorithm and its variations, and apply it in warranty context. We introduce a class of profit functions, which are functions of the warranty, and use it to formulate the problem of maximizing the company's profit over time as an online convex optimization problem. We use this formulation to present an approach to setting the warranty based on an online algorithm with low regret. Under a dynamic environment, this algorithm provides a warranty strategy for the company that maximises its profit over time.

*Key words:* warranty; online convex optimization; regret; profit functions

## 1 Introduction

A product warranty is an agreement offered by a producer to a consumer to replace or repair a faulty item, or to partially or fully reimburse the consumer in the event of a failure. Warranties are very widespread and serve many purposes, including protection for producer, seller, and consumer. They are used as signals of quality and as elements of marketing strate-

gies. A general treatment of warranty analysis is given by Blischke and Murthy [2, 3].

From the buyer's point of view, the main role of a warranty in any business transaction is protectional. Specifically, the warranty assures the buyer that faulty item will either be repaired or replaced at no cost or at a reduced cost. A second role of warranty is informational, as it implicitly sends out a message regarding the quality of the product and could influence buyer's purchase decision.

The main role of warranty from the producer's point of view is also protectional. Warranty terms may and often do specify the use and conditions of use for which the product is intended and provide for limited coverage or no coverage at all in the event of misuse of the product. A second important purpose of warranty for the seller is promotional. As buyers often infer a more reliable product when a long warranty is offered, this has been used as an effective advertising tool. In addition, warranty has become an instrument, similar to product performance and price, used in competition with other manufacturers in the marketplace.

Despite the fact that warranties are so commonly used, the study of warranties in many situations remains an open problem. This may seem surprising since the fulfillment of warranty claims may cost companies large amounts of money. Underestimating true warranty costs will result in losses for a company, overestimating them will result in uncompetitive product prices. The data relevant to the modeling of warranty costs in a particular industry are usually highly confidential, since they are commercially sensitive. Much warranty analysis therefore takes place in internal research divisions in large companies.

The common warranty parameters of interest analysed and evaluated are the expected warranty cost and the expected warranty cost per unit time over the warranty length for a particular item, as well as the life-cycle of the product, see Chukova and Hayakawa [4, 5]. Typically, the warranty length and the warranty policy are assumed to be known, which identifies the failure model. Based on the adopted model of the failure process, the total expected total warranty cost and, sometimes the variance of this cost, are evaluated.

The study presented here deviates from the traditional framework in warranty analysis. For simplicity we assume that the warranty is one-dimensional and non-renewing, i.e., the warranty is identified by its length and it starts at the time the item is sold or the service has began. We consider time periods, such that the manufacturer's profit functions, as functions of warranty, and the warranty may vary from time period to time period. In general, we assume that the optimal warranty and the profit functions are unknown, but the profit for the assigned warranty in any particular time period is known. The aim of this study, is to present an approach that will assure that if the warranty varies in a particular way, suggested

by an online algorithm, under reasonable, quite general assumptions on the profit functions, the long run average of the manufacturer's profit will be comparable with its profit if the optimal warranty were known at the time the product was launched on the market.

The outline of this paper is as follows. In Section 2 we present a brief overview of the online algorithm approach. The profit model is introduced in Section 3, and it is analysed, using an online algorithm, in Section 4. Section 5 contains concluding remarks.

## 2 Miscellaneous Results: Online Algorithm

In this paper, we concern ourselves with profit maximization, thus we consider the online convex programming problem with a sequence of concave functions and a maximization objective. In its simplest form, an online convex programming problem  $(F, \{p_1, p_2 \dots\})$  consists of a feasible region  $F \in \mathfrak{R}^n$  and an infinite sequence of concave "profit" functions  $\{p_1, p_2 \dots\}$ , each going from  $F$  to  $\mathfrak{R}$ . An algorithm for the online convex programming problem  $\mathcal{A}(F, \{p_1, p_2 \dots\})$  is an algorithm that produces a point  $w_i$ , which is a function only of the points  $w_1, w_2, \dots w_{i-1}$ , each previously produced by the algorithm, and the first  $(i - 1)$  functions  $p_1, \dots p_{i-1}$ . The regret of an algorithm is defined as

$$R(T) = \left( \max_{w^* \in F} \sum_{i=1}^T p_i(w^*) \right) - \sum_{i=1}^T p_i(w_i).$$

Interpreting, the regret measures the performance of the algorithm, which does not know  $p_i$  before producing  $w_i$ , to picking the single best point  $w^*$  in the feasible region  $F$  given knowledge of all the  $p_i$ 's in advance.

Online convex optimization, introduced by Zinkevich [8] was originally motivated by the notion of playing repeated games. Imagine playing an infinitely repeating game that proceeds in rounds. In round  $i$  we must pick a strategy knowing only the strategies we have chosen in the previous rounds, from 1 to  $(i - 1)$ , and the payoffs we received in those rounds. That is the motivation of the algorithm, which produces the point  $w_i$  knowing only  $w_1, w_2, \dots w_{i-1}$  and the first  $(i - 1)$  functions  $p_1, \dots p_{i-1}$ . Each  $w_i$  can be thought of as the strategy in the  $i^{th}$  round and each  $p_i$  can be thought of as the payoff function in the  $i^{th}$  round. The payoff function may change from round to round arbitrarily, since we do not know the strategies adopted by opponents in the game. In the repeated game settings, the regret then measures the amount of utility lost by a player who follows the strategy as specified by the algorithm versus picking the single best strategy to follow in all rounds.

Zinkevich exhibits an algorithm  $\mathcal{A}_f(F, \{p_1, p_2 \dots\})$ , in full-information

settings (see the Appendix), with regret  $R(T) = O(\sqrt{T})$ , which gives

$$\lim_{T \rightarrow \infty} R(T)/T \leq 0.$$

Interpreting, in the limit, following the strategies specified by the algorithm produces the same per period profit as picking the optimal single strategy. The quantity  $R(T)/T$  is commonly referred to as the average regret.

Online convex optimization has clear industrial applications. For example, consider a company producing a product. The company's profit could be a concave function of the warranty offered by the company. However, the profit does not only depend on the warranty, it could also depend on the types of products offered by competitors or the changing demands of customers. The profit function of the company in period  $i$  could be thought of as the  $p_i$  in the online convex optimization problem, and the  $w_i$  could be the warranty offered by the company in period  $i$ . An algorithm with low regret gives a warranty strategy for the company to follow that maximizes the company's profit over time.

One of the main hurdles to applying Zinkevich's algorithm directly is that it requires full knowledge of the function  $p_i$  after round  $i$ . In specific, Zinkevich's algorithm uses the gradient of the function  $p_i$ . However, in a realistic settings, such as the example in the previous paragraph, a company may not know the entire function  $p_i$ . Instead, all the company learns in round  $i$  is the value of  $p_i(w_i)$ . In other words, all the company learns is the amount of profit the company made in round  $i$ , not the entire profit function. Flaxman, Kalai, and McMahan [7] exhibit an algorithm for online convex optimization,  $\mathcal{A}_b(F, \{p_1, p_2 \dots\})$  in bandit settings (see the Appendix), using only the value  $p_i(w_i)$  of the profit function of the previous round and with regret  $R(T) = O(T^{\frac{3}{4}})$ .

Another concern with the direct application of online convex optimization is that the average regret results are in the limit as the number of rounds goes to infinity. Traditional industries, such as car manufacturing, have warranty on the order of years. Thus, even a few periods of the repeated profit maximization may take a human lifetime. However, warranties come in many varieties and today's markets can be largely autonomous. For example, consider a competition between online brokerage firms. A firm could offer a warranty on the amount of time required to execute a purchase or sell an order. The warranty offered could change dynamically throughout the trading day. The broker's customers could themselves be automated programs that dynamically choose which brokerage firm to use to execute trades. In such a scenario it is easy to imagine thousands of profit maximization rounds per day. Regardless of the plausibility of using online convex optimization in a specific application, the average regret results imply the startling conclusion that a company can attain nearly maximum profit in a dynamically changing environment,

without knowing anything about the future.

In this paper, we study online convex optimization as applied to the warranty applications described in this section.

### 3 The Profit Model

In what follows, we propose a general form of profit functions  $\{p_1, p_2 \dots\}$ , to be used with the two online convex optimization algorithms  $\mathcal{A}_f$  and  $\mathcal{A}_b$  for the warranty optimization examples in Section 2. Firstly, similarly to Bell, Keeney and Little [1], we define the market share function  $m(w)$  as a function of warranty  $w$  as follows

$$m(w) = \frac{a + g w}{a + g w + c}, \quad (1)$$

where  $a$  is a parameter of initial “attractiveness” or “reputation” of the company,  $g$  is the increase of the total attractiveness ( $a + g w$ ) of the company per unit increase of warranty, and  $c$  is the total attractiveness of the competitors of the company in the marketplace. It is easy to see that  $m(w)$  is an increasing function of  $w$ .

This form of the profit function is appropriate in modelling different market structures. For example, if  $c = 0$ , the company has a monopoly in the marketplace, whereas altering the value of  $c$  will model the arrival or departure of a competitor.

To gain some intuition on the market share function, suppose that the warranty  $w$  is zero. We then have

$$m(0) = \frac{a}{a + c}. \quad (2)$$

One can think of this equation as follows. Suppose a customer picks which company to use randomly, but with weights proportional to the company’s attractiveness. The form of  $m(0)$  in (2) is the probability the customer selects to do business with our company instead of a competitor, given that the company assigns no warranty to its products. Another interpretation of (2) is that, if the company assigns no warranty to its products it will have  $m(0)$  share of the market. Now, consider form (1) and let  $w \rightarrow \infty$ . We have  $\lim_{w \rightarrow \infty} m(w) = 1$ , which means that if the company offers a large warranty it will dominate the entire market.

Now, using the market share function  $m(w)$  given in (1), we introduce the profit function  $p(w)$ , again as a function of warranty. We propose:

$$p(w) = \mathbf{P} m(w) - \mathbf{R} m(w) F(w), \quad (3)$$

where  $\mathbf{P}$  is a constant equal to the total market value of the considered industry,  $\mathbf{R}$  is a constant equal to the penalty of total recall of all sold products and  $F(w)$  is the cumulative distribution function of the lifetime  $X$  of

the product. The latter function represents the quality and reliability of the production and governs the process of failures and related warranty claims. We assume a linear relationship between  $\mathbf{P}$  and  $\mathbf{R}$  of the following form

$$\mathbf{R} = \gamma \mathbf{P}, \quad \text{and} \quad 1.0 \leq \gamma \leq 2.0$$

In the case of  $\gamma \leq 1$ , even if all products are recalled, offering a large warranty will guarantee that the company will end up with a profit. On the other hand, if  $2.0 \leq \gamma$ , in order to avoid heavy penalties, the most appropriate strategy for the company is to sell the product with no warranty. Therefore, in both of these cases the optimal strategy of the company is known and we will focus our study to the nontrivial case of  $1.0 \leq \gamma \leq 2.0$ .

## 4 Modeling a Dynamic Environment

In what follows we display the performance of  $\mathcal{A}_f$  and  $\mathcal{A}_b$  in several differing models of a dynamic environment. First, we present an environment with a quality improvement under two failure scenarios: a gradual failure modeled with an exponential lifetime distribution, and a shock failure modeled with a Weibull lifetime distribution. Second, we present an environment with increasing competition again under two scenarios: a gradual increase in competition, and a shock increase in competition. Finally, we present an environment where we increase the penalty for faulty products. We show that in all these environments, the algorithms  $\mathcal{A}_f$ , and  $\mathcal{A}_b$  perform well as compared to algorithm “opt\_fixed” which selects a single, optimal warranty for all rounds, even though neither  $\mathcal{A}_f$  nor  $\mathcal{A}_b$  know the future profit functions. As algorithm  $\mathcal{A}_b$  is a randomized algorithm and its theoretical guarantees are in expectation, in each scenario we present the expected behavior of  $\mathcal{A}_b$  over 50 independent runs. In addition, we include in the comparison the algorithm “opt\_round” that selects the optimal warranty in each round.

### 4.1 Environment with Quality Improvement

Refer to the profit functions defined in (3). Our next goal is to use these functions for decision making related to warranty, in environment with quality improvement. We model the dynamic environment with quality improvement by using the cumulative distribution function  $F_X(w)$  of the lifetime  $X$  of the product. We consider two cases.

*Case 1.* Firstly, we assume that  $X \sim Exp(\lambda)$ , i.e.,

$$F_X(w) = 1 - e^{-\lambda w}, \quad (4)$$

and the mean time to failure is  $E(X) = \frac{1}{\lambda}$ . Based on (4) we define a sequence of profit functions

$$p_i(w) = \mathbf{P} m(w) - \mathbf{R} m(w)(1 - e^{-\lambda_i w}). \quad (5)$$

and use them in full-information settings, i.e., with  $\mathcal{A}_f(F, \{p_1, p_2 \dots\})$  as well as in bandit settings, i.e., with  $\mathcal{A}_b(F, \{p_1, p_2 \dots\})$ . The results are presented in Figure 1.

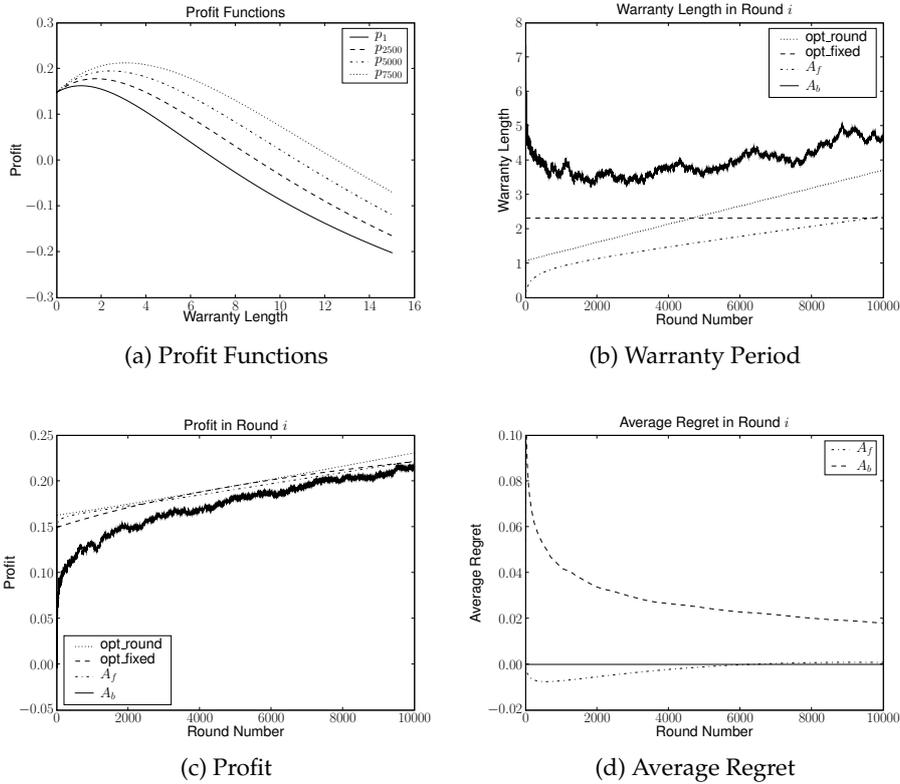


Figure 1: In Figure 1a are represented several profit functions, with  $p_i$  representing the profit function in round  $i$ . In Figure 1b are represented the warranty offered, as a function of the round number, by various algorithms. Figure 1c represents the profit earned, as a function of the round number, by each algorithm. Finally, Figure 1d represents the average regret of the two algorithms,  $\mathcal{A}_f$  and  $\mathcal{A}_b$ . The algorithm “opt.fixed” selects a single, optimal warranty for all rounds. The algorithm “opt.round” selects the optimal warranty in each round.

In this example, we model quality improvement by additively increasing the parameter of the exponential distribution representing the lifetime of the product. The mean of the distribution changes linearly from 4 to

8. The resulting profit functions are presented in Figure 1a. As you can see, in later rounds, as the quality increases, the company can offer a larger warranty to capture a larger fraction of the market and thus receive higher profit. Figure 1b shows the warranty offered by the various algorithms. Algorithm  $\mathcal{A}_f$  starts by offering a zero warranty, the imposed initial starting point, and follows an upward trend as the rounds increase. Algorithm  $\mathcal{A}_b$  has an initial starting point, dictated by the algorithm itself, around the middle of the feasible region. In all our examples, the feasible region is  $\{w \mid 0 \leq w \leq 15\}$ , i.e., the acceptable warranty is between 0 and 15. That is why, initially, the warranty of  $\mathcal{A}_b$  decreases from 7.5 and then increases as the rounds increase. Figure 1c shows the profit earned in each round by each algorithm. The figure illustrates the benefits of using  $\mathcal{A}_b$ , as it closely follows the profit received by optimizing the warranty in each round, but it assumes very limited information of the profit functions. Figure 1d shows how the average regret of  $\mathcal{A}_b$  decreases to zero as the rounds increase. In other words, the per period loss of  $\mathcal{A}_b$  as compared to following the optimal fixed warranty decreases to zero as the rounds increase. Even better results are pictured in Figure 1d for  $\mathcal{A}_f$ , however, it assumes knowledge of the gradient of the profit function in each round, whereas  $\mathcal{A}_b$  only assumes knowledge of the evaluation of the profit at a single point.

*Case 2.* Secondly, we assume that  $X \sim Weibull(\gamma)$ , i.e.,

$$F_X(w) = 1 - e^{-w^\gamma}, \quad (6)$$

and the mean time to failure is  $E(X) = \Gamma(1 + \frac{1}{\gamma})$  and create the sequence of profit functions

$$p_i(w) = \mathbf{P} m(w) - \mathbf{R} m(w)(1 - e^{-w_i^\gamma}). \quad (7)$$

In this example, we introduce quality improvement with a Weibull lifetime distribution. In a Weibull distribution, there is a sharp threshold at which most products fail. That is why in Figure 2a, the profit functions fall sharply as the warranty increases. Figure 2c represents the profit of the various algorithms. Notice that the profit for “opt\_fixed” begins negatively and sharply increases as the rounds increase. This is because the single warranty chosen by “opt\_fixed” in early rounds is greater than the failure threshold of the product, but is less than the failure threshold in later rounds. The profit earned by  $\mathcal{A}_b$  is negative in early rounds, since  $\mathcal{A}_b$  begins with an initial point in the middle of the feasible region, which is much larger than the failure threshold of the initial Weibull distributions. As the rounds increase,  $\mathcal{A}_b$  decreases the warranty, as pictured in Figure 2b. Since both the failure threshold increases and  $\mathcal{A}_b$  decreases the warranty in later rounds,  $\mathcal{A}_b$  eventually begins to make a profit. In late rounds,  $\mathcal{A}_b$  begins to approach the performance of  $\mathcal{A}_f$  and “opt\_round,” which outperform

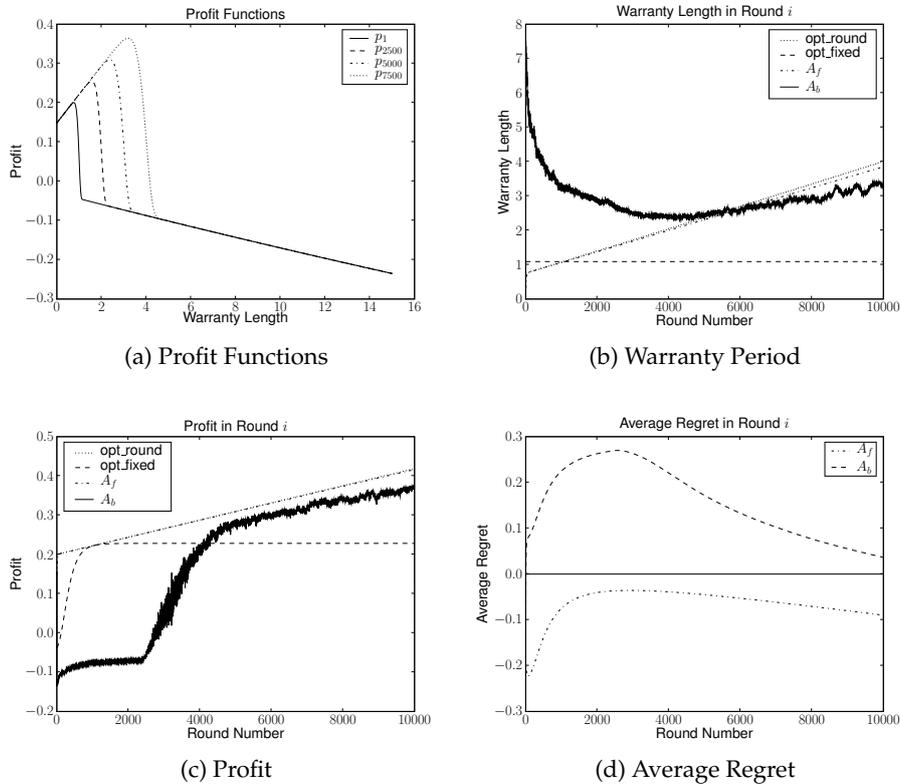


Figure 2: These figures represent the algorithms’ performance with a Weibull lifetime distribution. In a Weibull lifetime distribution, there is a sharp threshold before which most products are functioning properly and after which most products have failed. That is why in Figure 2a we see profit functions that fall sharply as the warranty increases.

“opt\_fixed,” since they can increase the warranty as the failure threshold increases. As expected, since  $A_f$  outperforms “opt fixed,” the average regret for  $A_f$ , pictured in Figure 2d, is negative. The average regret for  $A_b$  increases in the early rounds, while  $A_b$  is making poor profit, and quickly decreases in the later rounds.

## 4.2 Environment with Increasing Competition

We model the increase in competition in the profit function through the parameter  $c$  included in the market share function (1). In this example, we additively increase the competition from 2 to 50, with the parameter  $a$  set to 1. Interpreting, this means that initially the company has roughly 1 to 2 odds of attracting a customer. Towards the final round, the company has only 1 to 50 odds of attracting a customer, thus the competition has

increased. Figure 3b, through the graph for “opt\_round,” shows that the warranty that should be offered by the company increases as competition increases. This is to capture a larger fraction of the market as dictated by expression (1). The warranty of  $\mathcal{A}_b$  decreases throughout, as it initially begins at the middle of the feasible region. Algorithm  $\mathcal{A}_f$ , on the other hand, begins initially with a warranty of zero and closely follows the performance of “opt\_round.” Figure 3d shows that  $\mathcal{A}_b$  loses less than 15% of the total profit at the end of the example. This percentage would decrease to zero as the rounds go to infinity, by the results of Flaxman et. al. [7].

This example shows the algorithm behavior when there is a shock increase in competition. In round 2000, the arrival of a competitor decreases the market share of the company significantly, the value of  $c$  jumps from 2 to 50. This leads to the sudden drop pictured in Figure 4c, for all warranty settings, even for “opt\_round.” Though the profit for all algorithms has a sudden drop, it is interesting to see the algorithm’s reaction in changing the warranty, pictured in Figure 4b. Again, due to the different information settings of  $\mathcal{A}_f$  and  $\mathcal{A}_b$ , algorithm  $\mathcal{A}_f$  is near the optimal setting before the competition increase and needs a short time to re-adjust after the increase. On the other hand  $\mathcal{A}_b$  begins, as usual, in the middle of the feasible region and is decreasing the warranty towards the optimal setting in the initial rounds. After the competition increase,  $\mathcal{A}_b$  continues to decrease the warranty but at a slower pace. Even though the warranty offered by  $\mathcal{A}_b$  seems far from the warranty offered by the other algorithms, Figure 4d shows that its regret, as a percentage of the total profit gained by “opt\_fixed,” is once again steadily decreasing towards zero.

### 4.3 Environment with Changeable Penalties

In this example, we study a linear increase in the penalty from a faulty product. In specific, we alter the ratio  $\gamma$  between  $\mathbf{P}$  and  $\mathbf{R}$  in the profit function (3) from 1 to 2. A larger  $\gamma$  models a larger cost to replacing a failed item. As can be seen in Figure 5a, as the penalty for a faulty product increases, the optimal warranty goes to zero. Figure 5b shows how the warranty of  $\mathcal{A}_f$  starts at zero, increases until it passes the optimal warranty offered by “opt\_round,” and decreases back towards zero. It can also be seen that  $\mathcal{A}_b$  starts with a warranty of 7.5 and decreases towards zero. Figure 5c shows a similar performance of “opt\_round”, “opt\_fixed,” and  $\mathcal{A}_f$ . In that figure, it is clear that  $\mathcal{A}_b$  starts with a poor performance, but in the long run approaches the performance of “opt\_fixed.” The graph in Figure 5d can be explained through understanding the performance of algorithm  $\mathcal{A}_b$ , which is outlined in the Appendix.

In our penalty example, the optimal warranty approaches zero quickly. So, algorithm  $\mathcal{A}_b$  cannot set a warranty close to zero because of the algorithm’s projection to a subset of the feasible region. As the algorithm’s pa-

parameter  $\alpha$  approaches zero,  $\mathcal{A}_b$  can set a warranty closer and closer to zero. Thus, we can expect the regret shown in Figure 5d to decrease towards zero at a speed of  $O(\sqrt[6]{n})$ , matching that of the parameter  $\alpha$ .

## 5 Conclusions and Future Research Directions

In this paper we have presented a framework for analysis of warranty using an online convex optimization algorithm. We have introduced a class of profit functions that can be used to model a competitive market with warranties. We have shown that under incomplete information regarding the future changes in the environment, the decision maker could choose a warranty strategy that achieves a profit similar to the profit, that could have been generated by the unknown optimal warranty. In specific, we use the results of Zinkevich and Flaxman et. al. to exhibit strategies achieving near optimal profits, i.e. strategies with regret approaching zero in a long term. We exhibit several settings of changing environment and show that in each of these, the online algorithms can provide a reasonable support in warranty related decision making.

This study demonstrates that it is feasible for a company to maximize profit through adjusting warranty in a dynamic environment, without knowledge of the current or future market conditions. However, the algorithms presented here do have explicit limitations that should be noted before use in a real environment. First, as most optimization algorithms, the algorithms presented in the paper are guaranteed to work for convex objective functions. However, if the profit function of the company is not convex, it is possible for the algorithm to get stuck in a local optimum. Furthermore, as mentioned earlier, some products, such as cars, may not be appropriate for use with these algorithms because of the real-time length of a round, which is on the order of years. As demonstrated, specifically for the bandit algorithm, a large number of rounds are required to approach the optimal warranty period.

Furthermore, we are able to identify two possible directions for further research. One option is to focus on reducing the limitations of the used online algorithms. It would be interesting to see if these algorithms can be coupled with existing algorithms for avoiding local optima. For example, is it possible to pair the bandit algorithm with simulated annealing? What would such a pairing do to the regret guarantees of the original bandit algorithm? Would such a pairing deliver good performance in avoiding local optima? Another possible direction for further research is to try to apply our results to a real data, related to the performance of brokerage firms. Firstly, it will be challenging to find the appropriate set of real data. Moreover, it would be interesting to come up with a method for estimating the parameters of the profit function from real data, parameters such as the

total market size, the failure CDF, and the market share as a function of warranty period. Such an estimation would make it possible to investigate the application of these algorithms in a realistic situation.

## References

- [1] Bell, D., Keeney, R., and Little, J., (1975), Market Share Theorem, *Journal of Marketing Research*, 12:2, p.136-141
- [2] Blischke, W., and Murthy, D.N.P., (1993), *Warranty Cost Analysis*, Marcel Dekker, Inc.
- [3] Blischke, W., and Murthy, D.N.P., (1996), *Product Warranty Handbook*, Marcel Dekker, Inc.
- [4] Chukova,S., and Hayakawa,Y., (2004), Warranty Cost Analysis: Non-Zero Repair Time, *Applied Stochastic Models in Business and Industry*, 20 (1): pp. 59-71.
- [5] Chukova,S., and Hayakawa,Y., (2004), Warranty Cost Analysis: Renewing Warranty with Non- Zero Repair Time, *International Journal of Reliability, Quality and Safety Engineering* , 11(2), pp.1-20.
- [6] Hasan, E., Agarwal, A., and Kale, S., (2007), Logarithmic regret algorithm for online convex optimization, *Machine Learning* 96:169-192
- [7] Flaxman, A., Kalai, A., and McMahan, H.B., (2005), Online Convex Optimization in Bandit Settings: gradient descent without the gradient, In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, p.385 - 394.
- [8] Zinkevich, M., (2003), Online Convex Programming and generalized Infinitesimal Gradient Ascent, In *Proceedings of the twentieth international conference on machine learning (ICML)*, p.928-936.

## Appendix

### The On-line Algorithm

As mentioned earlier, we concern ourselves with profit maximization. Thus, consider an online convex programming problem consisting of a maximization objective, a feasible region  $F \in \mathbb{R}^n$  and an infinite sequence of concave functions  $\{p_1, p_2, \dots, p_i, \dots\}$ , each going from  $F$  to  $\mathbb{R}$ . We present the main ideas of the online algorithm in two different settings. Firstly, in full-information settings, when the profit function  $p_i$  is fully known after each round and secondly, in bandit settings, when the profit function  $p_i$  is

unknown, and only its value  $p_i(w_i)$  is revealed after the  $i^{\text{th}}$  round.

### Assumptions and Definitions

- The feasible region  $F \in \mathfrak{R}^n$  is
  - a bounded set, i.e., for any  $x, y \in F$ , there exists  $N \in \mathfrak{R}$ , so that  $d(x, y) \leq N$ , where  $d(x, y) = \|x - y\|$  and  $\|x\| = \sqrt{x \cdot x}$
  - a closed set, i.e., for any sequence  $\{w_i\}_1^\infty$ ,  $w_i \in F$ , if there exists  $x \in \mathfrak{R}^n$  such that  $x = \lim_{i \rightarrow \infty} w_i$ , then  $x \in F$ .
  - a non-empty set.
  - a convex set.
- The profit functions are differentiable.
- There exists  $N \in \mathfrak{R}$ , so that for all  $i$  and for all  $x \in F$ ,  $\|\nabla p_i(x)\| \leq N$ .
- For all  $y \in \mathfrak{R}^n$ , there exists an algorithm to produce  $\operatorname{argmin}_{x \in F} d(x, y)$ .
- For all  $i$ , there exists an algorithm, given  $x$ , to get  $\nabla p_i(x)$ .
- The projection of  $y$  over  $F$  is  $P(y) = \operatorname{argmin}_{x \in F} d(x, y)$ .
- The regret of  $\mathcal{A}$  until  $T$  is  $R_{\mathcal{A}}(T) = \left( \max_{x^* \in F} \sum_{i=1}^T p_i(x^*) \right) - \sum_{i=1}^T p_i(w_i)$ .
- A function  $p(x)$  satisfies an  $L$ -Lipschitz condition if there exists a real constant  $L$  such that  $d(p(x), p(y)) \leq L d(x, y)$ .

### Online Gradient Descent Algorithm $\mathcal{A}_f$ in full-information settings

Assume that the profit function  $p_i$  are fully known after the  $i^{\text{th}}$  round. Select an initial  $w_1 \in F$  and an updating sequence  $\eta = \{\eta_1, \eta_2, \dots, \eta_i, \dots\}$  with each  $\eta_i \in \mathfrak{R}^+$ . In time step  $(i+1)$ , after evaluating the profit function  $p_i(w_i)$ , move to the next point, which is

$$\mathcal{A}_f : x_{i+1} = P(x_i - \eta_i \nabla p_i(w_i)). \quad (8)$$

Assuming that the updating sequence  $\eta$  has the form  $\eta_i = \frac{1}{\sqrt{i}}$ , Zinkevich [8] has shown that the regret of the algorithm  $\mathcal{A}_f$ , given in (8) is

$$R_{\mathcal{A}_f}(T) \leq \frac{\|F\|^2 \sqrt{T}}{2} + \left( \sqrt{T} - \frac{1}{2} \right) \|\nabla p\|^2 ..$$

Therefore,

$$\limsup_{T \rightarrow \infty} \frac{R_{\mathcal{A}_f}(T)}{T} \leq 0,$$

where

$$\|F\| = \max_{x,y \in F} d(x,y), \quad \text{and} \quad \|\nabla p\| = \sup_{x \in F, i \in \{1,2,\dots\}} \|\nabla p_i(x)\|^2.$$

Imposing stronger assumptions on the profit functions and choosing appropriately the step sizes, Hasan, Agarwal and Kale [6] have extended Zinkevich's ideas by proposing several algorithms achieving logarithmic regret.

### Online Gradient Descent Algorithm $\mathcal{A}_b$ in bandit settings

In bandit settings, after the  $i^{\text{th}}$  round, the profit function  $p_i$  is unknown, and only its value  $p_i(w_i)$  is revealed. Therefore, the gradient of  $p_i$ , needed for  $\mathcal{A}_f$ , can not be accessed directly. The main difficulties in bandit setting is to obtain a one-point estimate of the gradient  $\nabla p_i(w_i)$ . Algorithm  $\mathcal{A}_b$  works as follows. It has a sequence of points  $y_i$  at which it would like to perform gradient descent, as in algorithm  $\mathcal{A}_f$ . However, to estimate the gradient at  $y_i$ ,  $\mathcal{A}_b$  selects a uniformly random point  $w_i$  from a small circle around  $y_i$ . Algorithm  $\mathcal{A}_b$  then sets  $y_{i+1}$  to  $y_i$  shifted in the direction of  $w_i$  with distance proportional to  $p_i(w_i)$ . To be sure that  $y_{i+1}$  is in the feasible region, the algorithm does a projection to a subset of the feasible region that has a small border around it. The reason for this projection to a subset is that future estimates of the gradient using randomly chosen points in a small circle should be entirely contained in the feasible region.

The algorithm then has three main parameters that change as the round number increases. Using the notation of Flaxman et. al, the first parameter  $\delta$  denotes the radius of the small circle around  $y_i$  from which we choose a uniformly random point  $w_i$ . The second parameter,  $\nu$  denotes the distance with which we move in the direction of the chosen point  $w_i$ . And the final parameter  $\alpha$  denotes the border that we keep around the subset of the feasible region. Each of these parameters goes to zero as the round number increases. The parameters  $\delta$ ,  $\nu$ , and  $\alpha$  go to zero at speeds of  $O(\sqrt[3]{n})$ ,  $O(\sqrt{n})$ , and  $O(\sqrt[6]{n})$ , respectively. Flaxman, Kalai, and McMahan [7] have shown that if the profit functions are  $L$ -Lipschitz the guarantee on the expected regret of  $\mathcal{A}_b$  is  $O(T^{\frac{3}{4}})$ . Moreover, if no Lipschitz or bounded gradient assumptions are placed on the profit functions, the guarantee on the expected regret is  $O(T^{\frac{5}{6}})$ . For more details, see Flaxman, Kalai, and McMahan [7].

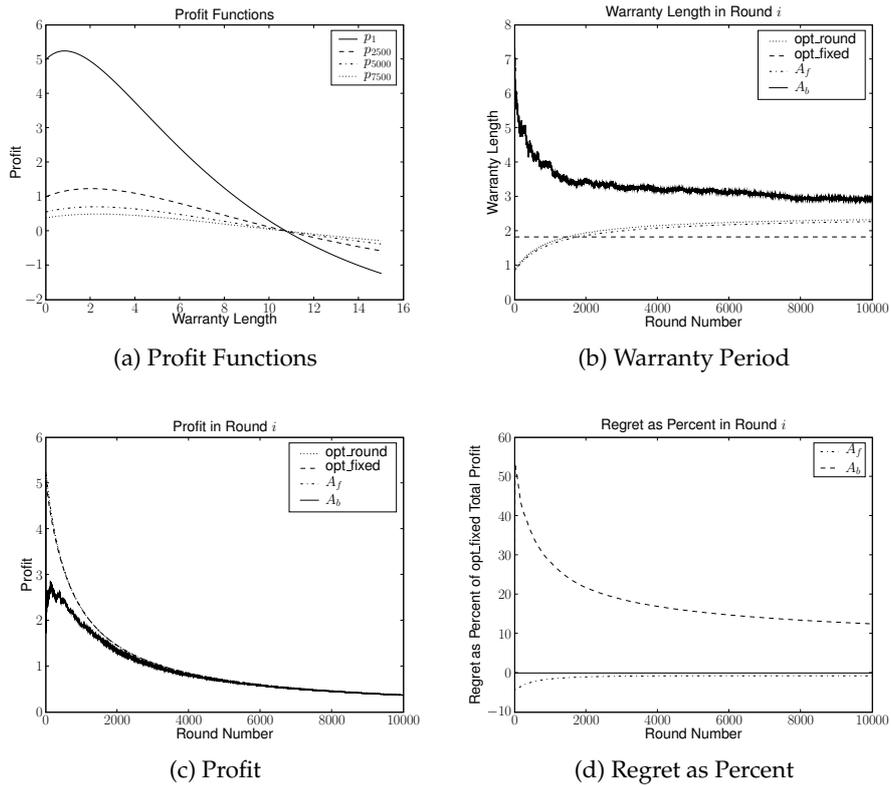
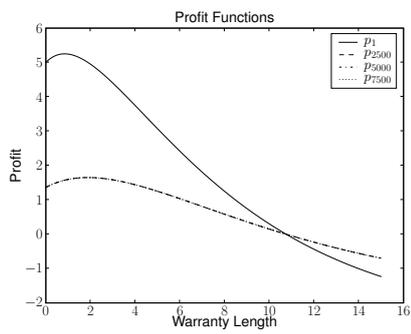
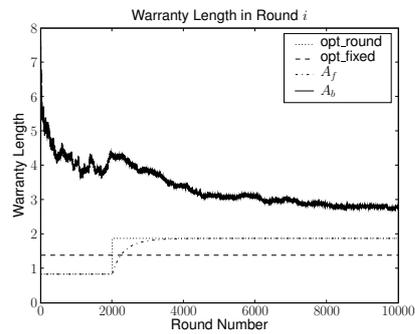


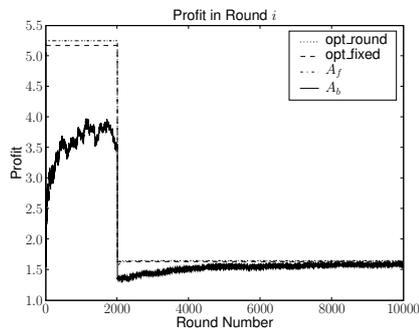
Figure 3: In this example, the competition increases additively from 2 to 50. As the competition increases, the company's share of the market decreases and so does the profit as shown in Figure 3a. Figure 3d, shows the regret in round  $i$  as a percentage of the total profit received by "opt.fixed" from rounds 0 through  $i$ . It illustrates that the theoretical results showing that the average regret tends to zero as the rounds tend to infinity translate into results showing that the percent of profit missed by  $A_b$  and  $A_f$  tends to zero as the round tends to infinity.



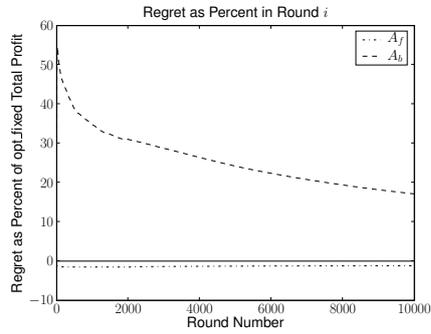
(a) Profit Functions



(b) Warranty Period

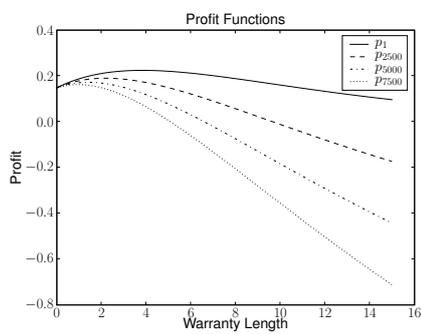


(c) Profit

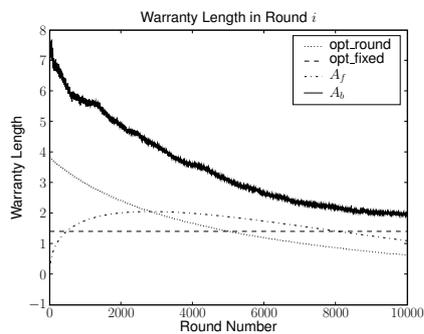


(d) Regret as Percent

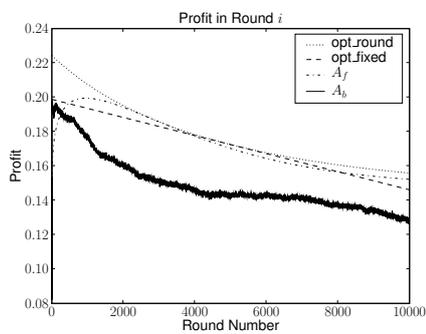
Figure 4: These graphs represent a shock increase in competition. In round 2000, the parameter  $c$  in the market share increases from 2 to 50.



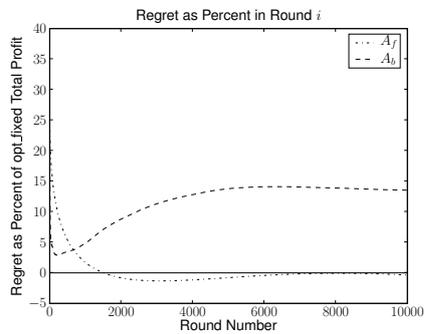
(a) Profit Functions



(b) Warranty Period



(c) Profit



(d) Regret as Percent

Figure 5: Additive increase in penalties