

Multi-Armed Bandit Strategies for Non-Stationary Reward Distributions and Delayed Feedback Processes

Larkin Liu
Richard Downe
Loblaw Digital
Toronto, ON, Canada

LARKIN.LIU@LOBLAW.CA
RDOWNE@LOBLAW.CA

Joshua Reid
University of Waterloo
Waterloo, ON, Canada

JS2REID@UWATERLOO.CA

Abstract

A survey is performed of various Multi-Armed Bandit (MAB) strategies in order to examine their performance in circumstances exhibiting non-stationary stochastic reward functions in conjunction with delayed feedback. We run several MAB simulations to simulate an online eCommerce platform for grocery pick up, optimizing for product availability. In this work, we evaluate several popular MAB strategies, such as ϵ -greedy, UCB1, and Thompson Sampling. We compare the respective performances of each MAB strategy in the context of regret minimization. We run the analysis in the scenario where the reward function is non-stationary. Furthermore, the process experiences delayed feedback, where the reward function is not immediately responsive to the arm played. We devise a new Bayesian technique (AG1) tailored for non-stationary reward functions in the delayed feedback scenario. The results of the simulation show superior performance in the context of regret minimization compared to traditional MAB strategies.

Keywords: Multi-Armed Bandit, Delayed Feedback, Non-stationary

1. Introduction

1.1 The Multi-Armed Bandit Strategy

In this paper, we present the formulation, algorithms and practical approaches to the application of different algorithms designed to effectively address the Multi-Armed Bandit (MAB) problem. We further investigate the practical case where the reward distributions of the arms are non-stationary along with the reward having a delayed feedback to the algorithm. This is more practical because in many cases, the time of retrieval of the reward information is not immediate. Furthermore, reward distributions may also change over time (non-stationary) in practical situations due to variance in the subjects being measured to determine the reward of each arm. We present simulated results in this paper to investigate the performance of different MAB strategies for both stationary and non-stationary reward functions. Specifically, we evaluate the performance of the MAB algorithms for the purpose of optimizing inventory availability for online eCommerce in the grocery sector.

The application of MAB strategies to the field of optimization has had a wide range of success. Applications of MAB strategies can be found in Web Search (Graepel et al., 2010), Real-Time-Bidding (Flajolet and Jaillet, 2017) and product pricing (Trovo et al., 2015). In this work, we specifically provide an overview and evaluation of MAB strategies for the purpose of forecasting item availability for online groceries. Each predictive algorithm provides a forecast of item availability for the entire inventory of items, and the MAB strategy will decide the optimal arm to select in order to optimize the empirical probability that each assorted article is actually available at the time of pick-up. It is essential to the operation of the business to be able to quickly identify the best algorithm for product assortment and play the arm that maximizes the metric to reduce the need for customers to substitute or remove items from their shopping list. In the following discourse, we shall study the application of MAB strategies in both the stationary and non-stationary scenario.

1.2 Industry Overview

In this paper, we address the issue of item availability via an eCommerce grocery service, created for the Loblaw Corporation, a large grocery store chain in Canada. In their online service, *Click-and-Collect*, customers can order items on the website, and proceed to collect the items in-store, or have the items delivered to their household via a third party. It is vital to the service that the assortment of items displayed on the website is readily available at the time of pick-up. Multiple algorithms run on the inventory and demand data to forecast the availability of these grocery articles. This paper, however, does not explicate the specific details of the in-house forecasting algorithms, rather, it addresses the application of MAB techniques to select the optimal algorithm to forecast the in-store assortment of products. We refer to these algorithms in this text simply as *assortment algorithms*, *models*, or *arms*. There are two major sections of this paper, first we present the theoretical background behind MAB algorithms. Subsequently we present the application of the MAB model in a simulated environment closely mimicking that of an actual online assortment process.

The primary metric, *fill-rate*, of each store is evaluated on a store basis and is the number of items that were in stock in that store when the order was being assembled. In other MAB experiments, multiple metrics are sometimes optimized by a single algorithm, some examples be found in (Öner et al., 2018). However, in our experiment we will consider the evaluation of each MAB strategy on a single metric, the *fill rate*.

2. Mathematical Background

2.1 Regret Minimization

We formulate the MAB strategy in the context of regret minimization. We define the Q function for every action k_t , as $Q(k_t)$, indicating that the reward obtained from playing arm, k , at time step t . Provided a policy π , the expected reward, V_t from taking action k_t can be expressed as,

$$V_t(\pi) = \sum^k Q(k_t)P(k_t|\pi) \tag{1}$$

k_t is the action undertaken in the current time, according to policy π . In our application, we defined V_t as the value of the state at time t . However, the reward function, $Q(k_t)$ is unknown. We denote the optimal value at time t as V_t^* , and under the optimal policy π^* .

$$V_t(\pi^*) = \max_{\pi} V_t(\pi) \tag{2}$$

The parameters, θ , which generate $Q(k_t)$ may or may not change with respect to time, depending on the stationarity of the reward function. Provided that any action k , can be selected depending on the policy, cumulative regret of a policy, upon completion of the game at the final time epoch T , is defined as R_T ,

$$R_T(\pi) = \sum_{t=0}^T [V_t(\pi^*) - V_t(\pi)] \tag{3}$$

Where $Q(k_t)$ is a constant for each k an time interval, t , depending on the stipulations in Section 2.2. The MAB strategy’s objective is to minimize the cumulative regret from measurements of the reward function of each arm played that epoch.

2.2 The Non-Stationary Bandit

We define the expected reward μ_t^k as the expected value of the reward from playing arm k at time t ,

$$\mu_t^k = \mathbb{E}[Q(k_t)] \tag{4}$$

A stationary bandit consists of a reward process where the expected reward of the reward distribution for each arm does not change over time. Hence for the stationary and non-stationary reward function μ^k and μ_t^k respectively,

$$\mu^k = f(\theta^k) \tag{5}$$

$$\mu_t^k = f(t, \theta^k) \tag{6}$$

Where θ^k represents the parameters of the expected reward distribution for arm k . In the non-stationary case, the reward function becomes a function of t as well as k . Where a deterministic function, f controls expected reward μ_t^k for purposes of simulation. Thus the optimal expected reward is calculated as,

$$\mu_t^* = \max_{k \in \{1, \dots, K\}} \mu_t^k \tag{7}$$

For example, (Besbes et al., 2014) investigated the effect of MAB’s in the non-stationary case using a sinusoidal function to generate μ_t^k , subsequently a confidence bound on the expected regret. The algorithm proposed in (Besbes et al., 2014) was unable to detect changes in the optimal arm dynamically, and relied on constantly resetting their MAB strategy to find the optimal arm in a brief time window. In our work, we propose an algorithm that is capable of detecting the change point of the reward distributions in a short amount of time, and is able to adaptively react to the change point of the expected reward function to further minimize regret.

2.3 The Delayed Feedback Framework

In our scenario, we define the concept of *delayed-feedback* as a reward process which does not immediately return an reward value upon playing of a selected arm. This idea is similarly noted in (Chapelle and Li, 2011), where some number of arms need to be played before a reward observation can be obtained from the system. In our case, we specify the number of stores as N , and the number of arms, or *assortment algorithms*, as K . Each arm is denoted as being the k^{th} arm, where $k \in \{1, \dots, K\}$. In our simulation, we specify $N = 100$, and $K = 10$.

In order to simulate our particular scenario, each store, out of N stores, is assigned an arm for the number of items to assort (or to fill). We refer to this number as γ . We set $\gamma = 50$ for all N stores. We estimate the parameter μ_t^k of the reward distributions at each delayed feedback interval. We express this estimate as,

$$\widehat{\mu}_t^k = \frac{1}{Nt} \sum_{t \in \Omega} \sum_{n=1}^N \frac{1}{\gamma} \sum_{i=1}^{\gamma} \mathbb{1}(k, \pi) Q(k_t) \tag{8}$$

Where $\mathbb{1}(k, \pi)$ is an indicator function stating if arm k was played under policy π , representing the MAB strategy. We use Ω to denote the set of all time epochs, t , contained within the observation period. In our delayed feedback scenario, once the store is assigned an arm, k , it must play that assortment algorithm for all γ items it needs to assort. Inherently, each arm, k , has a specified expected reward μ_t^k . In the stationary case, μ_t^k is constant and independent of time. In the non-stationary case μ_t^k can vary over time. For some strategies, we compute the parameter estimates for μ_t^k under a renewal framework, where only recent data is used to compute $\widehat{\mu}_t^k$, we refer to this as Ω_r , otherwise all observations from the beginning of the simulation is included in the time window, referred to as Ω ,

$$\Omega = \{1, \dots, t - 1\} \tag{9}$$

$$\Omega_r = \{t - r, \dots, t - 1\} \tag{10}$$

In the context of delayed feedback, the observations from the previous time epoch $t - 1$ until the end of the renewal period $t - r$ are used to calculate parameter estimates for $\widehat{\mu}_t^k$, as denoted by Eq. (10). Else if the MAB strategy is non-renewing, Eq. (9) is used to define the time window.

3. Multi-Armed Bandit Strategies

3.1 Epsilon Greedy Strategy

In the perspective of total cumulative regret, R_T , the ϵ -greedy selection strategy selects the expected optimal arm with probability $1 - \epsilon$, and other potential arm's with probability p_K . For n arms, we specify the probability of selecting the other arms as p_K is,

$$p_K = \frac{\epsilon}{K - 1} \tag{11}$$

In other words, each arm other than the optimal arm selected with probability p_n . For example, (Kuleshov and Precup, 2014) showed that the ϵ -greedy strategy performs significantly better than random experimentation. This is due to the fact the regret is greedily minimized as more observations are collected. However, the regret of the ϵ -greedy algorithm is constant at each iteration, and linear in t due to the fact that other sub-optimal arms will always be played. One can criticize an ϵ -greedy strategy on the grounds that it has poor asymptotic behavior, because it continues to explore long after the optimal solution becomes apparent.

Algorithm 1 ϵ -Greedy

```

1:  $Q = \emptyset$ 
2: for  $t = 0 \rightarrow T$  do
3:   for  $k = 1 \rightarrow K$  do
4:     Compute  $\mu_t^k$  from  $\Omega$ 
5:   end for
6:    $k = \operatorname{argmax}_k \widehat{\mu}_t^k$ 
7:   Play  $k$  with probability  $1 - \epsilon$ , else play another arm with probability  $p_K$ 
8:    $Q_k \leftarrow Q(k)$ 
9: end for

```

The lower bound on the expected regret of ϵ -greedy is proportional to T linearly in the infinite time horizon as the ϵ -greedy strategy is forced to play some amount of sub-optimal arms ϵ percent of the time for exploration purposes.

3.2 Adaptive Greedy (AG1) Strategy

In the context of delayed feedback the player must play N arms before updating their belief about the system contained within each time window Ω . In the the Adaptive Greedy (AG1) strategy, we opt to play the optimal arm based on the previous time epochs estimate of μ_t^k for $N\epsilon$ times, and leave ample room for exploration. We define a strategy to address the effect of non-stationarity in the reward function via an adaptive Bayesian method derived from the ϵ -greedy strategy. In this scenario, the previous time epoch’s parameter estimates are used to estimate the optimal arm k_t^* , at time t . We play the corresponding arm that maximizes the reward value based on the previous time window’s estimate. For each time epoch containing N arms plays, we define n_t^* as the number of times to play the optimal arm k_t^* .

$$n_t^* = \left\lfloor N(1 - \epsilon) \right\rfloor \tag{12}$$

And for non-optimal arms,

$$n_t = \left\lfloor N \frac{\epsilon}{K - 1} \right\rfloor \tag{13}$$

The AG1 strategy ensures that parameter for estimates for μ_t^k are constantly being updated to reflect changes due to non-stationarity. We use the previous time frames observation window Ω_r to estimate μ_t^k for the current time epoch.

Algorithm 2 AG1

```

1:  $Q = \emptyset$ 
2: for  $t = 1 \rightarrow T$  do
3:   for  $k = 1 \rightarrow K$  do
4:     Compute  $\mu_t^k$  from  $\Omega_r$ 
5:   end for
6:    $k = \operatorname{argmax}_k \widehat{\mu}_t^k$ 
7:   for  $1 \rightarrow n_t^*$  do
8:     Play  $k$ 
9:      $Q \leftarrow Q(k)$ 
10:  end for
11:  for  $k' = 1 \rightarrow K$  do
12:    if  $k' \neq k$  then
13:      for  $1 \rightarrow n_t$  do
14:        Play  $k'$ 
15:         $Q \leftarrow Q(k')$ 
16:      end for
17:    end if
18:  end for
19: end for

```

Similar to the ϵ -greedy algorithm the AG1 strategy forces exploration of potentially sub-optimal arms. This gives rise to a linear lower bound on expected total regret. However, it does provide a means to detect new changes in the expected value of the reward function, as the process solely relies on the last time epoch's, $t - 1$, estimate of the expected reward $\widehat{\mu}_t^k$. The AG1 strategy is an algorithm that modifies the ϵ -greedy's parameter estimation procedure such that it does not take into account historical value beyond that of the corresponding time window Ω_r .

3.3 UCB1 Algorithm

A more robust solution to the MAB problem is the UCB1 algorithm introduced first in (Lai and Robbins, 1985) and further analyzed in (Auer et al., 2002). The UCB1 algorithm places a penalty on the number of times each arm is selected. The UCB1 algorithm can be applied in both the stationary and non-stationary case, with a logarithmic bound on regret, as evidenced in (Garivier and Moulines, Oct. 2011). UCB1 demonstrates a logarithmic bound on regret, as opposed to a constant bound imposed by ϵ -greedy. However, the rate of convergence towards an optimal arm may be slower than other greedy strategies. We specify a UCB specific metric for the k^{th} arm as m_j ,

$$m_t^k = \mu_t^k + \sqrt{\frac{2 \log t}{n(k)}} \quad (14)$$

Where $n(k)$ refers to the number of time the k^{th} arm was played within the observable time window. The action we take at time k_t is expressed as the the arm that maximizes m_t^k ,

$$k_t = \underset{k}{\operatorname{argmax}} m_t^k \quad (15)$$

Algorithm 3 UCB1 Algorithm

```

1:  $Q = \emptyset$ 
2: for  $t = 0 \rightarrow T$  do
3:   for  $k = 1 \rightarrow K$  do
4:     Compute  $\mu_t^k$ 
5:   end for
6:   Play  $k_t = \underset{k}{\operatorname{argmax}} m_t^k$ 
7:    $Q \leftarrow Q(k)$ 
8: end for
    
```

The UCB1 algorithm imposes a logarithmic bound on the regret. For the stationary case this has been proven in the work of (Auer et al., 2002) expressed in Eq. (16).

$$\mathbb{E}[R_T(\pi)] \geq 8 \sum_{k: \mu_t^k < \mu_t^*} \left(\frac{\log n(k)}{\mu_t^k - \mu_t^*} \right) + \left(1 + \frac{\pi^2}{3} \right) \left(\sum_{k=1}^K \mu_t^k - \mu_t^* \right) \quad (16)$$

3.4 Thompson Sampling

(Thompson, 1933) introduced a Bayesian method framework for implementing MAB strategies which accounts for uncertain reward distributions. The advances in Bayesian methods have made it easier to apply randomized probability matching with almost any reward distribution (Scott). (Chapelle and Li, 2011) showed a highly competitive empirical result comparing Thompson Sampling with other MAB strategies. Thompson Sampling has some similarity with the ϵ -greedy algorithm, the key difference being that ϵ -greedy takes $\widehat{\mu}_t^k$ as the true estimate of the reward distribution and maximizes the expected reward with that assumption, whereas Thompson Sampling draws a random sample from an expected distribution parameters, θ , playing the arm a with the probability that it is the optimal arm based on model parameters, $\widehat{\theta}$. The probability of selecting arm k , is expressed as,

$$P(k) = \int \mathbb{I}[\mathbb{E}[Q(k)] = \max_a Q(k)] P(\theta|R) d\theta \quad (17)$$

Where the integral need not be computed explicitly, as it suffices to compute θ via a Bernoulli process, returning $\{0, 1\}$, as expressed in Eq. 18. Each arm is played $\gamma * n_t$ times, where n_t is the number of stores assigned to that arm at time t , and that is the estimate of θ for that trial. Thus θ can be expressed as the grand average of the probability of success across all time epochs, expressed in Eq. 8. We sample from the distribution $P(\widehat{\theta}|Q)$, provided historical rewards Q .

$$\widehat{\theta}_k = \widehat{\mu}_t^k \in [0, 1] \quad (18)$$

Algorithm 4 Thompson Sampling

```

1:  $Q = \emptyset$ 
2: for  $t = 0 \rightarrow T$  do
3:   Estimate( $\theta$ )
4:   for  $k \in \{1, \dots, K\}$  do
5:     Compute  $\widehat{\mu}_t^k$ 
6:     Sample  $Q(k_t) \sim \widehat{\theta}_k$ 
7:   end for
8:    $k_t = \operatorname{argmax}_k Q(k_t)$ 
9:    $Q \leftarrow Q(k_t)$ 
10: end for

```

The Thompson Sampling strategy typically achieves superior performance in terms of regret minimization in comparison with other MAB strategies. We present evidence of this in Section 4.0.1. A theoretical lower bound on regret is provided in (Agrawal and Goyal, 2012), providing further detail. Serving as one of the few papers exploring the theoretical regret bounds on Thompson Sampling, (Agrawal and Goyal, 2012) proved that the lower bound on expected regret for Thompson Sampling strategies is logarithmic. In Section 4.0.1, we will see empirical evidence of Thompson Sampling strategy being superior to other MAB strategies in the stationary reward scenario, in terms of regret minimization.

4. Experimental Results

In this section, we present the results of simulations pertaining to the cases of a stationary, and a non-stationary reward functions. The performance of the functions with respect to regret minimization is evaluated. In the stationary case, a fixed parameter μ^k is generated for each arm k . In the non-stationary case, a sinusoidal wave with respect a time is generated for μ_t^k .

4.0.1 STATIONARY SCENARIO

We compare the performance of three MAB strategies, ϵ -greedy, UCB1, and Thompson Sampling in the stationary case. For ϵ -greedy we set $\epsilon = 0.1$. A simulation environment similar to a live scenario is constructed. We provide k arms, where $k = 10$. Subsequently, the MAB strategy will determine which algorithm is optimal. With observations up until the end of the last time epoch $t - 1$. We specify 50 stores playing any of k arms over a period of 100 epochs.

We present simulation results mimicking performance of the in store *fill-rate* optimization algorithm for stationary reward distributions. In our simulation, the parameters are specified in Section 2.3 where $N = 50$, $K = 10$, and $\gamma = 50$. This allows for a total of 2500 arm plays each epoch. We measure the results in Table 2, on cumulative regret R_T and cumulative reward V_T . Note that the regret for an epoch can be lower than 0 due to the stochastic nature of the simulated assortment reward possibly producing rewards higher than the highest probability assigned to an assortment algorithm. We examine that Thompson Sampling obtains the best results in terms of cumulative regret, however, the

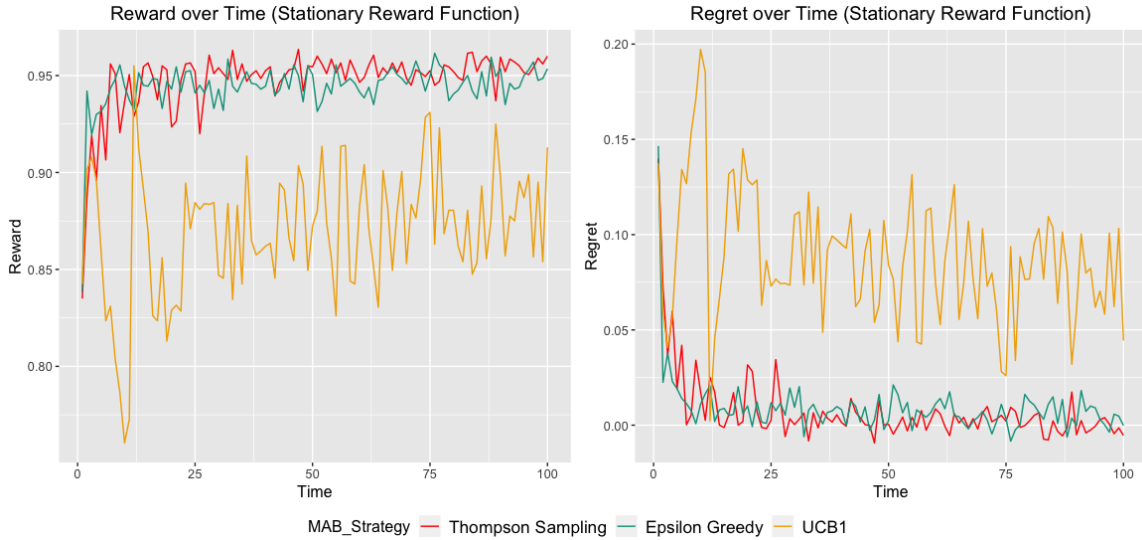


Figure 1: Regret over time for stationary reward function.

results are comparable to ϵ -greedy when measured against the cumulative reward. The simulation results indicate that the strategy results in poor performance compared to its counterpart strategies. We hypothesize this is due to the fact that the UCB1 strongly penalizes playing of the previous arm, causing it to over-explore less optimal strategies. Since in the delayed feedback scenario, arms must be played γ times in order to obtain a reward measurement, this hinders the performance of UCB1. Suffice it is fair to acknowledge that the UCB1 strategy has a logarithmic regret bound that is not constantly bounded. However, the over-exploration of sub-optimal arms in the delayed feedback scenario make the UCB1 strategy a poor choice for our application.

MAB Strategy	Cumulative Regret	Cumulative Reward
ϵ -Greedy	0.8928	94.45
Thompson Sampling	0.7262	94.75
UCB1	8.7477	86.98

Table 1: Table to test captions and labels

4.0.2 NON-STATIONARY SCENARIO

In this section we present the results of a MAB strategy for non-stationary reward functions. In one simulation we use two reward functions $K = 2$, where the expected reward function μ_t^k is sinusoidal with respect to time. The sinusoids are not identical to one another, providing a time period where one strategy is strongly dominant over another strategy. During this period, the MAB strategy should determine which arm is optimal, and play that arm most frequently. We compare the AG1 strategy outlined in Section 3.2 to the ϵ -greedy algorithm,

with a time window $\Omega_T = 3$ days. We observe that, in terms of regret minimization, the AG1 algorithm is superior to the ϵ -greedy algorithm. The reasoning for this is due to the fact that the AG1 strategy retains the parameter estimates from the previous time window, without the need to re-estimate the parameter again.




Figure 2: Regret over time for stationary reward function.

As evidenced from Fig. 2 for illustrative purposes, we see that the AG1 algorithm is able to minimize the regret, by changing k arm only when required to do so due to a regime change in the expected reward. The utilization of parameter memory, conducting a greedy strategy off of previous estimates of μ_{t-1}^k subsequently storing the data for future estimates provides a simple yet effective solution over an ϵ -greedy that simply restarts at each time interval.

MAB Strategy	Cumulative Regret	Cumulative Reward
ϵ^* -greedy	6.881	58.92
AG1	2.784	63.08

Table 2: Simulation results using $k = 2$ arms (sinusoidal reward function).

Subsequently, we investigate the results of this simulation for $K = 10$ arms. We include the results for ϵ^* -greedy and TS^* , where $*$ expresses the existence of a restart period. The restart period begins with no observations, playing each arm equally, until it re-discovers the optimal arm for the current time epoch, t . We see that TS^* is able to perform better than ϵ^* -greedy in terms of regret minimization, however, the AG1 algorithm performs significantly better for regret minimization due to its ability to remember the previous time window's, Ω_r , and provide a parameter estimate based on those observations.

MAB Strategy	Cumulative Regret	Cumulative Reward
ϵ^* -greedy	7.121	59.16
TS^*	6.241	59.82
AG1	2.558	63.71

Table 3: Simulation results using $k = 10$ arms (sinusoidal reward function).

5. Conclusion

In the context of MAB strategies, we have observed the efficacy of multiple strategies in the face of delayed feedback in the stationary and non-stationary cases for reward functions. In the stationary case, Thompson Sampling displays the most promising results in terms of cumulative total regret. The UCB1 displayed poor results for our non-stationary use case, however we acknowledge that there exists adaptive tuning methods for UCB1, that can be further explored to minimize regret. We hypothesize that this is due to the delayed feedback mechanism reducing exploration too strongly after some time epochs have passed, subsequently causing the strategy to select non-optimal arms. In the non-stationary case, we devise an MAB strategy based on ϵ -greedy. It performs significantly better than either Thompson Sampling or ϵ^* -greedy. The Bayesian approach yields significantly better performance as the model parameters are estimated in the prior time window, and retained for the next time epoch estimation.

For any number of arms, the AG1 strategy is shown to be effective at accounting for the non-stationarity of the reward function. In the case where a simulated sinusoidal wave was used to generate the reward function, regret was significantly minimized as the AG1 strategy suffered less cumulative regret than other MAB strategies with a restart period. In future work, it would be useful to compare the effectiveness of our model against other adaptive ϵ -greedy strategies, such as the one proposed in (Todic and Palm, 2011). Furthermore, it would be useful to consider tuning methods that can reduce the cumulative regret for cases regarding UCB1. The results of this simulation provide sound evidence that the AG1 provides superior performance for regret minimization in the MAB framework specifically in the scenario of delayed feedback, and non-stationary reward processes. Further areas of improvement can include a customization of Thompson Sampling to include a means to apply the Adaptive strategy for parameter estimation. We recognize the potential for MAB strategies to have a drastic impact in the optimization of experimentation for many industry applications, and due to the realization that non-stationary reward functions do occur, we conducted this simulation to obtain further insight into such processes.

References

Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In Shie Mannor, Nathan Srebro, and Robert C. Williamson, editors, *COLT*, volume 23 of *JMLR Proceedings*, pages 39.1–39.26. JMLR.org, 2012. URL <https://dblp.uni-trier.de/db/journals/jmlr/jmlrp23.html>.

- Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.*, 47(2-3):235–256, May 2002. ISSN 0885-6125. doi: 10.1023/A:1013689704352. URL <https://doi.org/10.1023/A:1013689704352>.
- Omar Besbes, Yonatan Gur, and Assaf J. Zeevi. Optimal exploration-exploitation in a multi-armed-bandit problem with non-stationary rewards. *CoRR*, abs/1405.3316, 2014. URL <http://arxiv.org/abs/1405.3316>.
- Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 2249–2257. Curran Associates, Inc., 2011. URL <http://papers.nips.cc/paper/4321-an-empirical-evaluation-of-thompson-sampling.pdf>.
- Arthur Flajolet and Patrick Jaillet. Real-time bidding with side information. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5162–5172. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7101-real-time-bidding-with-side-information.pdf>.
- Aurélien Garivier and Eric Moulines. On upper-confidence bound policies for non-stationary bandit problems. In *Algorithmic Learning Theory*, pages 174–188, Oct. 2011.
- Thore Graepel, Joaquin Quiñero Candela, Thomas Borchert, and Ralf Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. In *Proceedings of the 27th International Conference on International Conference on Machine Learning, ICML’10*, pages 13–20, USA, 2010. Omnipress. ISBN 978-1-60558-907-7. URL <http://dl.acm.org/citation.cfm?id=3104322.3104326>.
- Volodymyr Kuleshov and Doina Precup. Algorithms for multi-armed bandit problems. *CoRR*, abs/1402.6028, 2014. URL <http://arxiv.org/abs/1402.6028>.
- T.L Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.*, 6(1):4–22, March 1985. ISSN 0196-8858. doi: 10.1016/0196-8858(85)90002-8. URL [http://dx.doi.org/10.1016/0196-8858\(85\)90002-8](http://dx.doi.org/10.1016/0196-8858(85)90002-8).
- Doruk Öner, Altug Karakurt, Atilla Eryilmaz, and Cem Tekin. Combinatorial multi-objective multi-armed bandit problem. *CoRR*, abs/1803.04039, 2018. URL <http://arxiv.org/abs/1803.04039>.
- Steven L. Scott. A modern bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, 26(6):639–658. doi: 10.1002/asmb.874. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/asmb.874>.
- William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. 25(3/4):285–294, December 1933. ISSN 0006-3444 (print), 1464-3510 (electronic). doi: <https://doi.org/10.2307/2332286>. URL <http://www.jstor.org/stable/2332286>.

Michel Tokic and Günther Palm. Value-difference based exploration: Adaptive control between epsilon-greedy and softmax. In *KI 2011: Advances in Artificial Intelligence, 34th Annual German Conference on AI, Berlin, Germany, October 4-7, 2011. Proceedings*, pages 335–346, 2011. doi: 10.1007/978-3-642-24455-1_33. URL https://doi.org/10.1007/978-3-642-24455-1_33.

Francesco Trovo, Stefano Paladino, and Nicola Gatti. Multi armed bandit for pricing. 2015.