Blacklisting the Blacklist in Online Advertising

Improving Delivery by Bidding for What You Can Win

Yeming Shi
Dstillery
470 Park Ave S
New York, NY 10016
yshi@dstillery.com

Ori Stitelman
Dstillery
470 Park Ave S
New York, NY 10016
ostitelman@dstillery.com

Claudia Perlich
Dstillery
470 Park Ave S
New York, NY 10016
claudia@dstillery.com

ABSTRACT

Every day, billions of online advertising slots are bought and sold through real time bidding (RTB). In RTB, publishers sometimes reject bids to deliver ads (impressions) for some brands, due to, for example, direct deals with other brands. Publishers rarely disclose which brands they blacklist to ad buyers. Buyers bidding for a blacklisted brand waste computing resources in a low latency environment and lose an opportunity to show a good ad for a different brand. Here we describe a dynamic system developed at Dstillery that detects these (publisher, brand) combinations based on ad auction win rates and limits bidding for them to the minimum. This system demonstrates 1) a significant increase in the win rates of our bids, 2) a sizable reduction of system load, and 3) effectiveness in finding qualified non-blacklisted brands to replace blacklisted brands to show ads for. The system allows us to deliver more ad impressions while making fewer bids. In addition, we develop and demonstrate a methodology of choosing the optimal exploration-exploitation balance of the problem.

CCS CONCEPTS

• **Information systems → Online advertising**; *Online auctions*;

KEYWORDS

Online Advertising, Real-Time Bidding

ACM Reference format:

Yeming Shi, Ori Stitelman, and Claudia Perlich. 2017. Blacklisting the Blacklist in Online Advertising. In *Proceedings of ADKDD'17, Halifax, NS, Canada, August 14, 2017,* 6 pages.
https://doi.org/10.1145/3124749.3124753

1 INTRODUCTION

Online display advertising monetizes approximately 50 billion dollars each year. Every day, billions of ad slots are bought and sold through real time bidding (RTB) on ad exchanges. In RTB, buyers such as Dstillery bid to show an ad on a publisher website (inventory). The ad exchange receives bids from multiple buyers and runs an auction to determine who wins the opportunity to serve an ad to the audience. The resulting ad showing by the winning buyer is known as an impression. Buyers represent marketers (consumer brands or agencies) and are rewarded for ad campaign performance (more product purchases or service subscriptions) and scale (more impressions served). For each received bid request (BRQ), a buyer needs to decide which marketer to bid for this opportunity and the bid price. This is achieved by the buyer's bidder, a highly sophisticated system that makes these decisions and bids in real time. The bidder chooses a marketer based on this opportunity's estimated values across multiple marketers' ad campaigns. Here value means the probability that the audience will click the ad (click-through rate CTR) or take an action (conversion rate CVR), such as purchasing the marketer's product or subscribing to the marketer's service shortly after viewing the ad. These probabilities are estimated from data using supervised learning algorithms[3]. Typical features of the models include the audience's web browsing history and the opportunity's context such as the webpage[8], geographic location, etc. The design of Dstillery's model building and ad targeting system can be found in [10]. The bidder selects the marketer for which this opportunity is the most valuable and calculates a bid price based on this value. Bid price is an increasing function of the estimated CTR or CVR, and the relation can be linear[8] or non-linear[13].

It is commonly believed that the highest bid wins the auction on exchanges, but there are exceptions. Publishers have blacklists of marketers for which they refuse to serve impressions, even in cases where the highest bid for an opportunity is made for such marketers[1]. This can happen, for example, when the publisher has a direct deal outside RTB with one brand (such as Audi) and commits to not showing ads for competitor brands (other luxury car brands). Publishers rarely notify buyers which marketers are in their blacklists. Buyers bidding for a blacklisted marketer waste computing resources in a low latency environment and lose an opportunity to show a good ad for a different brand. At Dstillery, we've built a system to automatically detect the content of a publisher's blacklist ("banned marketers" or "BM") and limit bidding for BMs to the minimum. The detection is based on win rate (WR), the number of impressions won divided by the number of bids. For each BRQ, we retain a list of high value marketer candidates to choose from, so that if some of these marketers appear in the blacklist, we can bid for a different marketer. Our BM list is updated daily in accordance with changes to the publisher's blacklist. We keep a small fraction of BRQs unaffected by the filter (control group), so that we can still

[1]In this situation, the highest non-blacklisted bid could pay the second highest non-blacklisted bid in a second price auction.

ADKDD'17, August 14, 2017, Halifax, NS, Canada

ACM ISBN 978-1-4503-5194-2/17/08…$15.00
https://doi.org/10.1145/3124749.3124753

make a small number of bids on BMs to record their win rates. Our results show that, with an appropriately tuned criterion for which marketers enter the BM list, we 1) increase WR by more than 20%, 2) reduce system load considerably by making fewer bids, 3) often find qualified alternative marketers to replace blacklisted ones to show ads for, and 4) deliver more ad impressions.

This paper is organized as follows: Section 3 describes our bidding process, which is also the data collection process to train our BM list. Section 4 gives the algorithm/criterion for identifying BMs. Section 5 shows results on BM list training and validation. Section 6 shows production performance. Section 7 studies the optimal exploration-exploitation balance in this problem. The methodology and equations there are applicable to similar problems.

2 RELATED WORK

There exists plenty of literature on applying machine learning to problems in online advertising. Most papers focus on building models (for example, [3], [4], [6], [7], and [9]) to predict an opportunity's value (click-through/conversion probability) and calculating the optimal bid price (bid optimization) of a BRQ based on model prediction (for example, [8] and [13]). Other notable topics include pacing[5], i.e. the allocation of ad campaign budget over a period of time to achieve smooth delivery while optimizing for performance. In contrast, not much attention is devoted to win rates, despite its importance in various aspects of a buyer's ad bidding. Win rates naturally enter into consideration in pacing, as treated in [5]. Taking win rates into account when calculating bid prices leads to better campaign performance[13]. [1], [2], and [12] models market price, a quantity related to win rate. [14] models a campaign's historical win rate with a survival model for CTR estimation and bid optimization. [13] and [14] use continuous parametric functions of the bid price for win rate prediction in bid optimization. However, these win rate estimations don't include publisher as a feature, and we aren't aware of any literature pointing out the (exchange, publisher, marketer) dependency of win rates. Our work is novel in that it finds the interaction of publisher and marketer with regard to win rate and takes advantage of such insights by limiting bidding on marketers that are blacklisted by a given publisher.

3 BIDDING PROCESS

Dstillery makes bids on a number of exchanges. A down-sampled fraction of bids are fully logged[2]. The impressions are also logged. These logs allow us to track down every logged bid and find out whether we win or lose the auction. We apply an algorithm, detailed in Section 4, to this data collected a few days in the past, to identify a list of BMs on each (exchange, inventory) combination. Figure 1 shows how the BM list is applied in our bidder. After we receive a BRQ, we first do some pre-processing, such as identifying the inventory, device, and audience, and applying various filters such as a non-human activity (bot) filter[11]. Next, we find a list of ad candidates to consider for bidding. Each ad candidate represents a different marketer. Then we apply the BM filter to all but a small fraction of BRQs. When the BM filter is applied, marketers in the BM list of the (exchange, inventory) combination of the BRQ are eliminated from consideration, and the highest priced ad from the remaining ad candidates wins the internal auction and is submitted for bidding. A small random sample of BRQs skip the BM filter and proceed directly to the internal auction. For each of these BRQs, the highest valued marketer among all candidates is selected for bidding regardless of any blacklist.

[2]Because of the huge number of bids we make every day, logging every bid would be too costly. All bids presented in this paper refer to logged bids. Bid down-sampling does not affect WR.

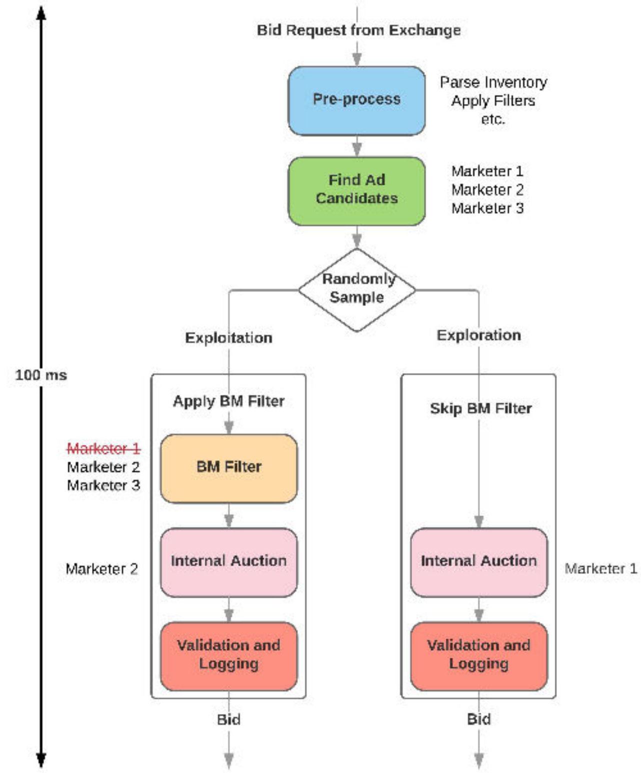

Figure 1: BM filter flow diagram. Marketers 1, 2, and 3 represent ad candidates in descending order of bid price. Marketer 1 happens to be in the BM list for this opportunity. If the BRQ skips the BM filter (right path), marketer 1 is selected for bidding. Otherwise (left path), marketer 2 is selected.

4 BANNED MARKETER IDENTIFICATION ALGORITHM

The intuition behind the algorithm is to use auction win rates on exchanges to uncover a publisher's blacklist. We analyze data on bids and impressions over the past $T = 8$ days. We choose $T = 8$ because marketers may enter and exit a publisher's blacklist (for example, because of the start and end of the publisher's direct deals), and $T = 8$ allows us to detect and react to any drop and increase in WRs within 8 days. We run the algorithm and refresh the list daily, and the list is applied in our bidder shortly after being refreshed.

4.1 General Approach

The WR of a blacklisted marketer's bids would be lower compared to other marketers' bids on the same exchange and publisher at the same price level. We cannot assume the WR of blacklisted marketers to be strictly zero because we aren't aware of how other machines make decisions. We need to consider bid price because otherwise we would mis-identify marketers as banned just because their bids have lower prices. This leads to the following algorithm.

First, discretize all bids into price buckets. For each (exchange e, inventory i, marketer m, price bucket p) combination, compute the expected WR as the aggregated win rate for all marketers m' on exchange e, inventory i and in price bucket p.

$$\mathbb{E}\left[WR_{e,i,m,p}\right] = \frac{\sum_{m'} Imps_{e,i,m',p}}{\sum_{m'} Bids_{e,i,m',p}} \quad . \tag{1}$$

Next, compute the resulting expected number of impressions

$$\mathbb{E}\left[Imps_{e,i,m,p}\right] = Bids_{e,i,m,p} \cdot \mathbb{E}\left[WR_{e,i,m,p}\right] \quad . \tag{2}$$

For a marketer blacklisted by a publisher, the actual number of *total* impressions for this (exchange, inventory, marketer) combination, aggregated over all price buckets,

$$TotImps_{e,i,m} = \sum_p Imps_{e,i,m,p} \tag{3}$$

would be much lower than[3] the expected number of *total* impressions

$$\mathbb{E}\left[TotImps_{e,i,m}\right] = \sum_p \mathbb{E}\left[Imps_{e,i,m,p}\right] \quad , \tag{4}$$

and that difference is used as the criterion for identifying a marketer as a BM on this (exchange, inventory).[4]

4.2 The Criterion and Parameter Tuning

$TotImps_{e,i,m}$ is a count. Under the null hypothesis, it follows a Poisson distribution. Both the mean and variance of $TotImps_{e,i,m}$ are equal to $\mathbb{E}\left[TotImps_{e,i,m}\right]$. Instead of computing the p-value of the Poisson distribution and declaring a BM based on p-value, we take a simpler approach. We identify a marketer to be a BM if[5]

$$TotImps_{e,i,m} < \alpha \cdot \mathbb{E}\left[TotImps_{e,i,m}\right] - \beta \cdot \sqrt{\mathbb{E}\left[TotImps_{e,i,m}\right]} \tag{5}$$

with appropriate constants α and β to be discussed soon. A choice of $\alpha = 1, \beta = 3$, for instance, would arise if we approximate the Poisson distribution with a Gaussian and use a 3 standard deviation rule, which corresponds to a 0.14% p-value. The greater $\mathbb{E}\left[TotImps_{e,i,m}\right]$, the more exact the Poisson to Gaussian approximation. With small $\mathbb{E}\left[TotImps_{e,i,m}\right]$, Equation (5) has the desired property that it won't identify a marketer as banned even if $TotImps_{e,i,m}$ is zero. For instance, with $\alpha = 1, \beta = 2$, $\mathbb{E}\left[TotImps_{e,i,m}\right]$ needs to be at least 5.9 in order for the RHS of Equation (5) to be positive. This is consistent with the requirement of the p-value of Poisson distribution being sufficiently small.

What values of α and β should be used? Experience proves that $\alpha = 1.0$ is not a good choice. What matters here is the business problem, not a purely statistical problem. We want to detect (exchange, inventory, marketer) combinations with very low (price-adjusted) WR (BM combinations) because of publishers' *blacklists*. $\mathbb{E}\left[TotImps_{e,i,m}\right]$ grows much faster than $\sqrt{\mathbb{E}\left[TotImps_{e,i,m}\right]}$ at large values of $\mathbb{E}\left[TotImps_{e,i,m}\right]$. Under $\alpha = 1.0$, if $\mathbb{E}\left[TotImps_{e,i,m}\right]$ is large, Equation (5) will be too aggressive in identifying a marketer as banned. For example, under $\alpha = 1.0$ and[6] $\beta = 3.0$, if $\mathbb{E}\left[TotImps_{e,i,m}\right] = 1000.0$, Equation (5) will become $TotImps_{e,i,m} < 905$. But we shouldn't put a combination in BM list if it is expected to win 1000 impressions[7] and actually wins 904. The WR of this combination being below average by only about 10% is irrelevant for our blacklisting. A good choice of α should be less than 1.0. In fact, we choose

$$\alpha = 0.6 \quad , \tag{6}$$

$$\beta = 1.2 \quad . \tag{7}$$

With this choice of parameters, the RHS Equation (5) is only positive if $\mathbb{E}\left[TotImps_{e,i,m}\right] \geq 7.0$, so a BM combination needs to have $\mathbb{E}\left[TotImps_{e,i,m}\right] \geq 7.0$ for it to be detected.

5 TRAINING AND VALIDATION

In the previous section, we described an algorithm to detect BMs. We'd like to verify that we have not overfitted the data and that the changes in win rates are not so fast that the learning is not effective in execution. In order to validate the algorithm in a way that reflects how we intend to execute it in production, we perform an out-of-time test, i.e. test the BM list on future data. Each BM list is generated based on data from the past eight days and used in our system for one single day. Therefore, for validation, we generate the BM list based on eight days' data (estimation set) and monitor WRs for (exchange, inventory, marketer) combinations in and outside the list on these eight days as well as on the following day (test set). Since the criterion for making the BM list compares actual with expected WRs, validation is also based on comparing actual against expected WRs. Table 1 shows, for both the eight-day estimation set and one-day test set, the aggregated actual WR vs. the aggregated expected WR for combinations in and outside the BM list[8]. The "Bids" column is the number of bids in each row.

---

[3]This algorithm is loosely motivated by the χ^2 statistical test. Under the null hypothesis that the marketer in question has the same WR as other marketers with the same (exchange, inventory, price bucket) combination, the actual number of impressions we win for this (exchange, inventory, marketer, price bucket) combination $Imps_{e,i,m,p}$ should have mean equal to $\mathbb{E}\left[Imps_{e,i,m,p}\right]$ and variance equal to $\mathbb{E}\left[Imps_{e,i,m,p}\right]$. Hence, the χ^2 statistic for all price buckets,

$$\chi^2_{e,i,m} = \sum_p \frac{(Imps_{e,i,m,p} - \mathbb{E}\left[Imps_{e,i,m,p}\right])^2}{\mathbb{E}\left[Imps_{e,i,m,p}\right]}$$

should follow a χ^2 distribution with degrees of freedom equal to the number of price buckets. A marketer having $Imps_{e,i,m,p}$'s lower than $\mathbb{E}\left[Imps_{e,i,m,p}\right]$'s and an extremely large χ^2 value would indicate that it has a significantly lower WR on this publisher compared to other marketers.

[4]In our bidding history, we see on average 30 distinct marketers on each (exchange, inventory) combination.

[5]The RHS of Equation (5) should be understood as having a floor sign.

[6]Under $\alpha = 1.0$, 3.0 is already a very large choice of β for identifying BMs at small $\mathbb{E}\left[TotImps_{e,i,m}\right]$. Under $\alpha = 1.0$ and $\beta = 3.0$, Equation (5) only starts to find BMs if $\mathbb{E}\left[TotImps_{e,i,m}\right] \geq 11.0$.

[7]Combinations with 1000 or more expected impressions in $T = 8$ days is common. They contribute more than 50% to our total impressions.

[8]For example, the aggregated actual WR for combinations in BM list L is

$$\frac{\sum_{(e,i,m)\in L} \sum_p Imps_{e,i,m,p}}{\sum_{(e,i,m)\in L} \sum_p Bids_{e,i,m,p}} \quad .$$

Expressions for expected WRs are similar, but with $\mathbb{E}\left[Imps_{e,i,m,p}\right]$.

Table 1: Actual win rate (WR) vs. expected win rate ($\mathbb{E}[WR]$) for (exchange, inventory, marketer) combinations in and outside the BM list.

Evaluation Set	in BM list	Bids	WR	$\mathbb{E}[WR]$
Estimation	False	$1.24e9$	5.99%	5.31%
	True	$2.31e8$	0.81%	4.49%
Test	False	$1.22e8$	6.37%	5.75%
	True	$2.34e7$	0.94%	4.12%

The table indicates that in the estimation set, combinations outside the BM list have slightly higher actual (5.99%) than expected WR (5.31%). The same is true in the test set (6.37% vs. 5.75%). Similarly, combinations in the BM list have much lower actual than expected WR in the estimation set (0.81% vs. 4.49%) as well as test set (0.94% vs. 4.12%). These demonstrate the success of our algorithm in identifying combinations with much lower WRs than expected and that such trends continue into the future. The algorithm identifies 30 thousand BM combinations, and not bidding on them can considerably reduce the load of our bidder. From the data in Table 1, we can estimate the improved system global WR as a result of applying the BM filter. In the test set, if we had only bid for combinations outside the BM list, the global WR would have been 6.37% on this particular day.

6 IN VIVO PERFORMANCE RESULTS

Having observed good out-of-time validation results, we applied the algorithm and BM filter (as illustrated in Figure 1) in production. The BM list/filter positively affects our system's performance in several aspects.

First, the win rate of the treatment group (the group for which the BM filter is applied) is significantly higher than the win rate of the control group (the group for which the BM filter is skipped). Table 2 shows a WR comparison between the two groups for the first two days that the BM filter is applied in our system[9]. "Lift" is the treatment group's WR divided by the control group's WR and is more than 1.20 on both days. During the 4-month period from filter deployment to the time of this writing, the BM filter has given daily WR lifts of 1.23 ± 0.03 . One can estimate the treatment group's WR using the method described at the end of Section 5. The treatment group's observed WR matches well with such estimation. In fact, on average the treatment group's observed WR is more than 95% of the estimated value. The small discrepancy exists because, when the top marketer candidate is filtered in the treatment group, we bid for an alternative marketer with a lower price on average compared to bidding on the same exchange and inventory and for the same alternative marketer in the control group[10]. Therefore, bids in the treatment group have a lower average WR compared to outside-BM-list bids in the control group.

[9]WRs in Tables 2 and 1 differ because they are WRs on different dates. WR depends on the bidding competition level in the market.

[10]As a result, on an average day, the treatment group's average bid price is lower than the control group's by 6%. The average price paid to show an impression, or CPM, is 2% (not 6% because of second-price ad auctions) lower in the treatment group.

Table 2: Global win rates of control and treatment groups for the first two days following BM filter deployment.

Day	Control Group	Treatment Group	Lift
Day 1	3.37%	4.17%	1.24
Day 2	3.23%	3.96%	1.23

Furthermore, not bidding on combinations in the BM list results in a 20% to 30% system load reduction in terms of the number of bids.

Last but not least, when a BM is filtered, there is a 48% chance that we find a qualified alternative marketer to bid in our system. This and the WR increase together let us deliver 6% more net impressions while making 14% fewer net bids (after accounting for bids made for the alternative). Finding an alternative marketer helps improve the alternative marketer's ad campaign performance (clicks or conversions). Each campaign has a delivery goal with respect to the number of impressions, which is smoothly allocated over the duration of the campaign. Every day, our pacing engine activates just the right number of highest ranked opportunities for each campaign to meet its daily allocation of delivery. In the absence of a BM filter, we lose active opportunities to show ads to qualified alternative marketers. Consequently, the pacing engine has to activate lower ranked opportunities to deliver for the alternative marketers, which leads to poorer campaign performance for them.

7 OPTIMAL EXPLORATION

In this section we present a methodology that is applicable to a class of similar problems. Our task is to find the optimal random sampling rate to allocate opportunities to the control group. There is a tradeoff between large and small sampling rates. As a general principle, the optimal point is where the number of times spent on exploring bad opportunities, in our case bidding for blacklisted combinations, is minimal. If the sampling rate is too large, we will spend too many bids on opportunities that are hard to win. On the other hand, if the sampling rate is too small, we will not make enough bids for some banned (exchange, inventory, marketer) combinations to make statistically significant judgements about their quality on a later date, which could result in us bidding on them later and bidding on them *more* on average over time. In the extreme case, if the control group sampling rate were zero, we would make no bids on BM combinations. Sooner or later, the past T days would have insufficient data for the algorithm to detect those BMs, and consequently we could bid heavily on them the next day. There is an optimal (possibly non-zero) control group sampling rate that minimizes the number of bids made to a BM combination. We need exploration[11] to help us make fewer bids for BMs.

[11]A well-known explore-exploit algorithm is UCB1 for the multi-armed bandit problem. Similar to our approach, UCB1 also explores bad opportunities. Our problem is different from multi-armed bandit in that 1) we want to block blacklisted marketers rather than pick the best option, 2) there are more variables in our problem such as bid price, and 3) a marketer's win rate can change over time. One may still take a UCB-like approach for our problem. The equations will be different from UCB1. Our control group sampling rate tuning to minimize the "regret" is novel, and our approach may have a lower regret than a UCB-like approach.

7.1 Number of Bids Made for a BM

Here we consider an (exchange, inventory, marketer) combination that has a very low win rate and should be in the BM list. We calculate the total number of bids *per day* made to this BM combination, in the control and treatment groups together. Our purpose is to find the control group sampling rate such that this BM combination receives a minimal number of bids. As we will see, the combination could be caught in the BM list on some days and escape detection on other days, resulting in us making fewer or more bids for it on different days. We will compute the average number of bids made for it per day.

In order to calculate this, we need to know on which days a BM combination is caught/not caught by our algorithm. Any criterion we choose to express

$$TotImps_{e,i,m} \ll \mathbb{E}\left[TotImps_{e,i,m}\right] \tag{8}$$

implicitly requires $\mathbb{E}\left[TotImps_{e,i,m}\right]$ to be sufficiently large. If $\mathbb{E}\left[TotImps_{e,i,m}\right]$ is too small, we can't declare it to be a BM even if $TotImps_{e,i,m}$ is zero. This is regardless of how we choose to express the condition in Equation (8). In particular, with Equation (5) as the criterion and the choice of parameters in Equations (6) and (7), $\mathbb{E}\left[TotImps_{e,i,m}\right] \geq 7.0$ in the past T days is required[12]. For $\mathbb{E}\left[TotImps_{e,i,m}\right]$ to be sufficiently large, a certain minimal number of bids (denoted by N) needs to be made for this BM combination in the past T days. In fact, N bids made to a BM should also be the sufficient condition[13] for its detection. Under criterion Equation (5) and the choice of parameters in Equations (6) and (7), N depends on the (exchange, inventory, marketer) combination, via

$$N \cdot \mathbb{E}\left[WR_{e,i,m}\right] = 7.0 \quad , \tag{9}$$

where $\mathbb{E}\left[WR_{e,i,m}\right]$ is this marketer's expected win rate on this exchange and inventory, aggregated over all price buckets. Expected WR in each price bucket is an average over all marketers (see Equation (1)) and doesn't depend on the specific marketer. Table 1 also indicates that the expected WR aggregated over price buckets doesn't depend strongly on whether the (exchange, inventory, marketer) combination is in the BM list. Therefore, we will approximate $\mathbb{E}\left[WR_{e,i,m}\right]$ with our global WR 4% for simplicity. Substituting this into Equation (9) gives $N = 175$. A BM is detected only if at least $N = 175$ bids is made for it in the past $T = 8$ days.

Now we're ready to compute the average number of bids per day made for a BM. It turns out that, at a given control group sampling rate, BM combinations belong to three classes depending on their sizes. Each class of BM has a distinct behavior with regard to being detected by our algorithm or not. BMs of small size are never caught by the algorithm, BMs of intermediate size enter and exit our BM list in a periodic pattern, and BMs of large size will remain in the BM list every day. Denote by c the control group random sampling rate. Consider a BM combination that gets $n_{e,i,m}$ bids per day in the absence of a BM filter. How often it enters the BM list depends on $n_{e,i,m}$. Obviously, if $n_{e,i,m} < N/T$, the combination will never enter the BM list, regardless of c. Under this circumstance, we will make $n_{e,i,m}$ bids for it every day. On the opposite end, a large BM combination with $n_{e,i,m} \geq N/(c \cdot T)$ will always be caught in the BM list. We will make $c \cdot n_{e,i,m}$ bids for such a BM every day. Lastly, BMs of intermediate size, i.e.

$$n_{e,i,m} \in \left[\, \frac{N}{T} \,,\, \frac{N}{c \cdot T} \,\right) \quad , \tag{10}$$

will enter and leave the BM list. For each of them, we'd like to compute the fraction of days it stays in the BM list. Assume without loss of generality that we turn on the BM filter on day 1. Since we've made $n_{e,i,m} \cdot T \geq N$ bids for it from day $-(T-1)$ to day 0, the algorithm catches this combination on day 1. Let x be the number of days it remains in the BM list. Because it is in the BM list on day x, the number of bids made to it in the T days before day x cannot be smaller than N, i.e.

$$n_{e,i,m} \cdot (T - x + 1) + c \cdot n_{e,i,m} \cdot (x-1) \geq N \quad . \tag{11}$$

Similarly, because it leaves the BM list on day $x+1$, we have

$$n_{e,i,m} \cdot (T - x) + c \cdot n_{e,i,m} \cdot x < N \quad . \tag{12}$$

Solving Equations (11) and (12) for integer x gives

$$x = 1 + \left\lfloor \frac{T - N/n_{e,i,m}}{1-c} \right\rfloor \quad . \tag{13}$$

In the T days before day $x+2$, the number of bids made to this BM combination is still the same as in the T days before day $x+1$ and thus less than N. Hence, the combination still won't be detected on day $x+2$. It is not hard to see that this combination will be outside the BM list from day $x+1$ to day $T+1$. It will re-enter the BM list on day $T+2$, because in the T days before day $T+2$ it is in the BM list for $(x-1)$ days. Day $T+2$ is the beginning of a new $(T+1)$-day cycle for this BM with regard to detection. The $(T+1)$-day cycle will keep on repeating itself. Therefore, in a $(T+1)$-day period, the combination is in the BM list for x days, and the average number of bids made per day for this BM combination is

$$b_I(n_{e,i,m}) = \frac{n_{e,i,m} \cdot (T + 1 - x) + c \cdot n_{e,i,m} \cdot x}{T+1} \quad , \tag{14}$$

where x is given by Equation (13). Putting these together, the number of bids per day $b_{e,i,m}$ made to any BM combination is

$$b_{e,i,m} = \begin{cases} n_{e,i,m} & \text{if } n_{e,i,m} < N/T \\ c \cdot n_{e,i,m} & \text{if } n_{e,i,m} \geq N/(c \cdot T) \\ b_I(n_{e,i,m}) & \text{otherwise} \end{cases} \quad , \tag{15}$$

where $b_I(n_{e,i,m})$ is given by Equation (14). The left plot in Figure 2 shows $b_{e,i,m}$ against c for $N = 175$, $T = 8$, and $n_{e,i,m} = 150$. As c is decreased down to zero (from right to left in the plot), the two discontinuities (jumps) in $b_{e,i,m}$ represent points where the BM combination transitions from staying in the BM list every day to leaving the list for 1 out of every $T+1$ days and then to leaving the list for 2 out of every $T+1$ days. The number of discontinuities in the plot depends on the parameters ($n_{e,i,m}$, etc). For this BM, $b_{e,i,m}$ is minimized at 2% (the left discontinuity in the plot). A different BM combination would have a different optimal c. For simplicity, we choose a single optimal c for all BMs. This is a restriction in our implementation, but it can change if we decide to be more sophisticated. To find this single optimal c, we shall consider all BMs.

[12] See the end of Section 4 for an explanation of this point.

[13] Since a BM has a significantly lower win rate than expected (see Table 1). If we make enough bids for it, it has a high chance of meeting the criterion Equation (8) and enter the BM list.

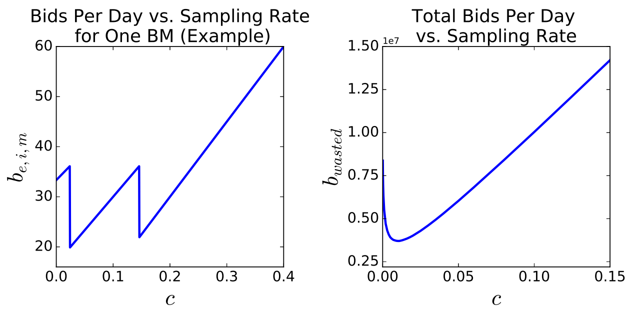

Figure 2: Left: an example of bids per day made for a BM as a function of control group sampling rate, as calculated by Equation (15). The parameters are $N = 175$, $T = 8$, and $n_{e,i,m} = 150$. Right: bids per day made to all BMs against control group sampling rate, as calculated by Equation (16).

7.2 Minimizing Bids Made for All BMs

Here we find all BM combinations and their $n_{e,i,m}$ numbers, compute the number of bids made for each combination, sum them up, and minimize the sum with respect to c . To identify all BM combinations, we look at data over a long period of time (20 days) and apply the criterion in Equation (5). We choose the period to be before we turn on the BM filter in our system, so that we have accurate $n_{e,i,m}$ numbers for the BM combinations. A period of 20 days enables us to collect enough data and detect practically all *true* banned combinations. Based on the list of true BMs L_t, along with the $n_{e,i,m}$ for each of them, we can calculate the $b_{e,i,m}$ for each of them using Equation (15) and sum up these $b_{e,i,m}$'s to get the number of bids per day make for all BMs, b_{wasted}.

$$b_{wasted} = \sum_{(e,i,m) \in L_t} b_{e,i,m} \quad . \tag{16}$$

We perform this calculation for a range of c's, and find the c that minimizes b_{wasted} . The right plot in Figure 2 shows the total bid count b_{wasted} versus different values of c. In our system, the optimal c turns out to be 1%. Setting the control group sampling rate above 10%, for example, is obviously unnecessary. Setting it to zero is undesirable, too. In either case, b_{wasted} can be significantly reduced by moving c towards its optimal value.

7.3 Generality of the Analysis

The method described in this section has general applications to a class of problems. In these problems, we want to select good opportunities, such as BRQs that we can win, and discard bad opportunities, such as BRQs that are hard to win. The quality of opportunities could change over time. We infer the quality from data and exploit the insights. In order to make statistically significant judgements, we need to explore each opportunity at least N times. This includes exploring known bad opportunities with a probability c. This section presents a way to find the optimal value of c. Assume we decide on the quality of opportunities based on data in the past T days[14]. The number of explorations per day for a bad opportunity is given by Equation (15), where n is the number of times per day we naturally "see" this opportunity, and subscript (e, i, m) is understood as an index of the bad opportunity. These equations are completely general. To find an optimal c for a bad opportunity, one minimizes Equation (15) with respect to c.

8 CONCLUSIONS

In this paper we have detailed an algorithm for identifying consumer brands that certain websites refuse to show ads for on each ad exchange. Our algorithm improves online ad bidding win rate by 23%, reduces system load, increases delivery, and improves performance. In addition, we developed and demonstrated a methodology of choosing the optimal balance between dedicating opportunities to exploration and dedicating opportunities to the execution of insights, where we block bad opportunities. If not enough opportunities are dedicated to exploration, one will end up taking more bad opportunities, due to a lack of estimation precision. The methodology is applicable to a class of problems where the goal is to identify and block bad opportunities and to react to possible changes in opportunity quality over time. The equations we obtained are general and based on minimal and reasonable assumptions.

REFERENCES

[1] O. Chapelle. 2015. Offline evaluation of response prediction in online advertising auctions. In *Proc. of ACM WWW*.
[2] Y. Cui, R. Zhang, W. Li, and J. Mao. 2011. Bid landscape forecasting in online ad exchange marketplace. In *Proc. of ACM SIGKDD*.
[3] B. Dalessandro, D. Chen, T. Raeder, C. Perlich, M. Han Williams, and F. Provost. 2014. Scalable hands-free transfer learning for online advertising. In *Proc. of ACM SIGKDD*.
[4] X. He, J. Pan, O. Jin, T. Xu, B. Liu, T. Xu, Y. Shi, A. Atallah, R. Herbrich, S. Bowers, and J. Q. Candela. 2014. Practical Lessons from Predicting Clicks on Ads at Facebook. In *Proc. of AdKDD*.
[5] K.-C. Lee, A. Jalali, and A. Dasdan. 2013. Real time bid optimization with smooth budget delivery in online advertising. In *Proc. of AdKDD*.
[6] K.-C. Lee, B. Orten, A. Dasdan, and W. Li. 2012. Estimating Conversion Rate in Display Advertising from Past Performance Data. In *Proc. of ACM SIGKDD*.
[7] A. K. Menon, K.-P. Chitrapura, S. Garg, D. Agarwal, and N. Kota. 2011. Response Prediction Using Collaborative Filtering with Hierarchies and Side-information. In *Proc. of ACM SIGKDD*.
[8] C. Perlich, B. Dalessandro, R. Hook, O. Stitelman, T. Raeder, and F. Provost. 2012. Bid optimizing and inventory scoring in targeted online advertising. In *Proc. of ACM SIGKDD*.
[9] T. Raeder, C. Perlich, B. Dalessandro, O. Stitelman, and F. Provost. 2013. Scalable supervised dimensionality reduction using clustering. In *Proc. of ACM SIGKDD*.
[10] T. Raeder, O. Stitelman, B. Dalessandro, C. Perlich, and F. Provost. 2012. Design principles of massive, robust prediction systems. In *Proc. of ACM SIGKDD*.
[11] O. Stitelman, C. Perlich, B. Dalessandro, R. Hook, T. Raeder, and F. Provost. 2013. Using co-visitation networks for detecting large scale online display advertising exchange fraud. In *Proc. of ACM SIGKDD*.
[12] W. C.-H. Wu, M.-Y. Yeh, and M.-S. Chen. 2015. Predicting winning price in real time bidding with censored data. In *Proc. of ACM SIGKDD*.
[13] W. Zhang, S. Yuan, and J. Wang. 2014. Optimal real-time bidding for display advertising. In *Proc. of ACM SIGKDD*.
[14] W. Zhang, T. Zhou, J. Wang, and J. Xu. 2016. Bid-aware gradient descent for unbiased learning with censored data in display advertising. In *Proc. of ACM SIGKDD*.

[14] Day is the unit of time in our application. Other applications can have different units of time, such as hour.