

Managing Risk of Bidding in Display Advertising

Haifeng Zhang[†], Weinan Zhang^{†,‡}, Yifei Rong[‡], Kan Ren[‡], Wenxin Li[†], Jun Wang[‡]

[†]Peking University, [‡]MediaGamma Ltd. & University College London

[‡]YOYI Inc., [‡]Shanghai Jiao Tong University

pkuzhf@pku.edu.cn, jun.wang@mediagamma.com, yifei.rong@yoyi.com.cn

ABSTRACT

In this paper, we deal with the uncertainty of bidding for display advertising. Similar to the financial market trading, real-time bidding (RTB) based display advertising employs an auction mechanism to automate the impression level media buying; and running a campaign is no different than an investment of acquiring new customers in return for obtaining additional converted sales. Thus, how to optimally bid on an ad impression to drive the profit and return-on-investment becomes essential. However, the large randomness of the user behaviors and the cost uncertainty caused by the auction competition may result in a significant risk from the campaign performance estimation. In this paper, we explicitly model the uncertainty of user click-through rate estimation and auction competition to capture the risk. We borrow an idea from finance and derive the value at risk for each ad display opportunity. Our formulation results in two risk-aware bidding strategies that penalize risky ad impressions and focus more on the ones with higher expected return and lower risk. The empirical study on real-world data demonstrates the effectiveness of our proposed risk-aware bidding strategies: yielding profit gains of 15.4% in offline experiments and up to 17.5% in an online A/B test on a commercial RTB platform over the widely applied bidding strategies.

Keywords

Risk-aware Bidding Strategy, Value at Risk, Demand-Side Platform, Real-Time Bidding, Display Advertising

1. INTRODUCTION

Display advertising has become a significant battlefield for big data. As the advertising transactions are aggregated across websites in real time, the display advertising industry has a unique opportunity to understand the internet traffic, user behaviors, and online transactions. In this paper, we have used datasets from a leading DSP (Demand Side Platform) in China, iPinYou, which processes up to 18 billion ad impressions per day [3]. The advertising platform ANONYM¹, which has deployed our proposed algorithms in this paper, currently handles 5 billion ad transactions daily. Moreover, Fikisu DSP claims to process 32 billion ad impressions daily [2]; Turn reports to handle 2.5 million per second in the peak time [29]. To get the picture of the scale, the New York Stock Exchange trades around 12 billion shares daily [1], while the Shanghai Stock Exchange trades about 14 billion shares daily [4]. It is fair to say that the transaction volume from display advertising has already surpassed that of the financial market.

¹Anonymized for blind review.

To further compare, similar to the financial market, ad impression trades are also automated largely by auction mechanisms — while the financial market uses a double auction to create bid and ask quotes [6], RTB (Real-time Bidding) display advertising adopts the second-price auction to gather bid quotes from advertisers once an impression is being generated [37, 23]. As it is possible now to track user actions resulted from an online campaign, advertising optimization becomes more resembling to that of the financial market trading and tends to be driven by the marketing profit and return-on-investment (ROI). That is, there is explicit and measurable campaign goal of acquiring new users from the campaign in order to obtain additional sales from the acquired users. Thus, how to properly bid an ad impression to drive the profit and ROI becomes essential to performance-driven campaigns [25, 38].

Bidding strategies are normally built by estimating the utility of each ad impression, which is commonly done by predicting the underlying user’s click-through rate (CTR) or conversion rate (CVR) [17]. Existing solutions for predicting CTR range from linear models [13, 21, 15, 17, 26] and gradient boosting decision trees (GBDT) [31] to factorization machines [24]. All of them aim to return a predicted CTR value of the given ad impression. However, we never know whether such a *point* estimation is confident enough. The true underlying CTR may heavily deviate from the predicted value, given the significant uncertainty of the underlying user behavior. Such utility fluctuation, along with the cost uncertainty caused by the auction competitions [5], results in a significant risk of the campaign profit estimation which should not be ignored.

In this paper, we depart from conventional CTR point estimations and explicitly model the CTR distribution to capture the uncertainty (risk) of the utility measure for each potential ad impression. Our idea is inspired from finance about *value at risk*: a statistical technique used to measure and quantify the level of financial risk with an investment [19]. The value at risk of an ad display opportunity is generally defined as the lower bound of the ad impression value guaranteed by a specific probability (risk). With the measured CTR prediction and market price risks, we propose two practical risk-aware bidding strategies to handle the uncertainty from both utility estimation and market competition. Our methods are evaluated on two large-scale real-world ad datasets. With campaign profit as the key performance indicator (KPI), we find that our risk-averse bidding strategies, which penalize the bids with high uncertain CTR (or profit) and award the confident ones, yield a 15.4% profit gain over a linear bidding strategy with a conventional logistic regression CTR estimator [25] and largely outperform the risk-neutral and risk-seeking bidding.

Furthermore, we have deployed the risk-averse bidding strategies on ANONYM’s DSP and conducted a 7-day on-

line A/B test. In the live test, our proposed risk-aware bidding strategies bring 17.5% higher campaign profit and 61.5% higher CTR over the conventional ones by effectively saving money on uncertain and low-value opportunities, which verifies the practical effectiveness of our solutions of managing advertising bidding risk.

The rest of this paper is organized as follows. We discuss the related work in Section 2. Our CTR distribution modeling is provided in Section 3. In Section 4, the concept of value at risk and the derived risk-aware bidding strategies are discussed. Offline and online experimental results are provided in Sections 5 and 6, respectively. We finally conclude this paper and discuss our future work in Section 7.

2. RELATED WORK

User Response Prediction. Predicting the probability of a specific user response, e.g., CTR and CVR, is a key function for performance-driven online advertising [13, 21, 15]. The applied CTR estimation models today are mostly linear. Logistic regression is the most widely used model, normally trained by stochastic gradient descent (SGD) [17, 26]. The authors in [21] proposed a novel online learning algorithm called follow-the-regularized-leader (FTRL) to train logistic regression from the streaming data. The model successfully bypasses the learning rate update problem in SGD and it empirically works effectively. Bayesian probit regression [13] is another linear model for online learning where the feature weights are modeled with a distribution and the model learning is via updating the weight posterior. Binary naive Bayes [14] is also a popular linear model, by assuming the features are conditionally independent.

Linear models are simple and effective in learning, but may fail to capture the interactions between the assumed (conditionally) independent raw features [13]. By contrast, non-linear models are capable of learning feature interactions in various ways and could potentially improve prediction performance [24, 31]. Gradient boosting decision trees (GBDT) [31, 15] are a straightforward non-linear model to capture feature interactions. Moreover, latent factor models, particularly factorization machines (FMs) [24], map each binary feature into a low dimensional continuous space, and the feature interaction is automatically explored via vector inner product.

Real-Time Bidding Strategies. The emergence of ad exchanges for display advertising in 2009 [23] provides automatic trading mechanism for advertisers to buy media inventory in impression level and determine the acceptable price via second price auction [37].

The authors in [12] proposed an algorithm that made dynamic bidding decisions to achieve an optimal delivery with the budget constraint. In [8], the bid price from each campaign was adjusted by the publisher or the supply-side platform in real time and the goal was to maximize the publisher side profit. Borrowing the idea of the optimal truth-telling bidding in sponsored search [10], a basic bidding strategy is to bid the estimated true value for each ad impression. For performance-driven campaigns, the predefined true value is normally based on the economical value of the user actions, such as a click or a conversion. The expected true value for a specific impression is estimated as action value multiplied by action rate [17, 8]. However, the truth-telling bidding strategy is optimal only when the budget and auction volume are not considered. With the campaign budget and lifetime auction volume constraints, the optimal bidding strategy may not be truth-telling. Extending from the truth-telling bidding strategy, the authors

in [25] proposed the generalized bidding function with a linear relationship to the predicted CTR for each ad impression being auctioned. Compared to [25], the authors in [39] proposed a functional optimization framework to directly optimize the bidding function, where the derived function showed that an optimal bidding function could be non-linear w.r.t. predicted CTR. The non-linearity is closely related to the distribution of market price [5]. Extending to multiple campaign bid optimization task, the authors in [38] further devised linear and non-linear bidding functions from the functional optimization framework. Extending from previous strategies, our work introduces a new factor, i.e. the standard deviation of predicted CTR, to be considered by bidding function in addition to predicted CTR.

Risk Management and Applications. Risk is a consequence of action taken in spite of uncertainty [22]. The objective of risk management is to assure uncertainty does not to-some-extent deflect the business from its goals [9].

Modern portfolio theory (MPT) [20] originates from modeling uncertainty of the return of combinations of multiple financial assets. It presents a quantitative method to measure such uncertainty (or risk) and embeds it into the decision making of investment [16]. In MPT, the variance of the return of each asset is modeled as its risk. Then the risk and expected return of a portfolio of invested assets are quantified by the fund allocation, the mean return of the assets and their covariance matrix [20, 16]. MPT utilizes the mean-variance analysis to make an investment portfolio for any tradeoff between the risk and the expected return, or w.r.t. a reference investment such as bank rates [28]. With such advantages, MPT has been adopted in almost everywhere of financial investment [11].

Recently, the ideas of risk management has been introduced to information retrieval, such as document ranking in web search [33, 34] and diversification in top-N recommendation [30, 41], to improve the model robustness or catch the users' satisfaction on uncertainty psychologically.

Computational advertising is associated with a certain level of deficit risk, particularly for performance-driven campaign as the goal is to acquire new users and gain more sales from them. The risk comes from the dynamics of the market and the user online behaviors [32]. The authors in [38] proposed to measure campaign-level risk and return in a special case of arbitrage between CPM and CPA. Compared to [38], our work focuses on single campaign optimization, and our risk is modeled from the uncertainty of user response and market competition at impression-level. Generally, our work borrows the concept of value at risk from finance to derive risk-aware bidding strategies intending to reasonably allocate budget between uncertain impressions and confident impressions and achieve a campaign-level profit gain, which is unlike in finance where risk control is only for balancing return and risk at item-level (impression-level in RTB).

3. CTR DISTRIBUTION MODELING

In this section, we explicitly model the CTR distribution in order to deal with the uncertainty of a CTR estimation. Next, in Section 4 we shall propose risk-aware bidding strategies from the inferred CTR distribution and the market price distribution.

3.1 Preliminary: Bayesian Logistic Regression

We propose to use a Bayesian logistic regression to model the CTR distribution due to the following reasons: (i) logistic regression (LR) has been widely deployed as the CTR

prediction model in most RTB ad platforms [17, 25] and our model is a natural extension to tackle the uncertainty of a CTR estimation with LR; (ii) We adopt Bayesian treatment to model uncertainty since it has been well studied by previous works [40, 13] for CTR estimation; (iii) Although Bayesian probit regression [40, 13] has the potential to model the uncertainty, its probit activation function is of no closed form, and thus is computationally low cost effectiveness in RTB.

For readability, in this section, we present a preliminary on Bayesian logistic regression, while for details, we refer to [7]. For a multi-dimensional feature vector \mathbf{x} representing the input ad display opportunity, the conventional logistic regression estimates the CTR by:

$$\hat{y} = \sigma(\mathbf{w}^T \mathbf{x}) = \frac{1}{1 + e^{-\mathbf{w}^T \mathbf{x}}}, \quad (1)$$

where σ is the sigmoid function and \mathbf{w} is the weight vector of logistic regression. The likelihood of observing the correct binary click label y given features \mathbf{x} and weights \mathbf{w} is

$$p(y|\mathbf{x}, \mathbf{w}) = \sigma(\mathbf{w}^T \mathbf{x})^y (1 - \sigma(\mathbf{w}^T \mathbf{x}))^{(1-y)}. \quad (2)$$

In the Bayesian version of logistic regression, \mathbf{w} is modeled as a random variable with a p.d.f. $p(\mathbf{w})$. Thus the marginal conditional probability $p(y|\mathbf{x})$ is

$$p(y|\mathbf{x}) = \int_{\mathbf{w}} p(y|\mathbf{x}, \mathbf{w}) p(\mathbf{w}) d\mathbf{w}. \quad (3)$$

We follow [40, 13] to adopt a Gaussian prior $N(\boldsymbol{\mu}_0, q_0^{-1} \mathbf{I})$ on \mathbf{w} , which is a practical setting. After observing a data instance (\mathbf{x}, y) , the posterior distribution of \mathbf{w} becomes

$$p(\mathbf{w}|\mathbf{x}, y) = \frac{p(y|\mathbf{x}, \mathbf{w}) p(\mathbf{w})}{\int_{\mathbf{w}'} p(y|\mathbf{x}, \mathbf{w}') p(\mathbf{w}') d\mathbf{w}'} \quad (4)$$

$$\propto \sigma(\mathbf{w}^T \mathbf{x})^y (1 - \sigma(\mathbf{w}^T \mathbf{x}))^{(1-y)} \prod_i \sqrt{\frac{q_{\text{old},i}}{2\pi}} e^{-\frac{(w_i - \mu_{\text{old},i})^2 q_{\text{old},i}}{2}},$$

where $\mu_{\text{old},i}$ and $q_{\text{old},i}$ are the prior parameters of the i -th dimension of \mathbf{w} before observing the data instance (\mathbf{x}, y) .

Approximation of Posterior. Since the posterior Eq. (4) is intractable, we maintain a Laplace approximation [7] to keep it consistent with the prior. There are alternative approximate inferences such as variational inference (VI) [7]. In this paper we adopt Laplace approximation due to its simpler implementation and lower computational cost than VI, which involves extra variational parameters and invokes a time-consuming EM algorithm for training. The Gaussian approximation to the posterior distribution takes the form

$$q(\mathbf{w}) = N(\mathbf{w}|\boldsymbol{\mu}_{\text{new}}, \mathbf{S}_{\text{new}}), \quad (5)$$

where $\boldsymbol{\mu}_{\text{new}}$ is defined by the \mathbf{w}_{MAP} which maximizes the logarithmic posterior:

$$\boldsymbol{\mu}_{\text{new}} = \arg \max_{\mathbf{w}} y \ln \sigma(\mathbf{w}^T \mathbf{x}) + (1 - y) \ln(1 - \sigma(\mathbf{w}^T \mathbf{x})) - \frac{1}{2} \sum_i q_{\text{old},i} (w_i - \mu_{\text{old},i})^2 + \text{const}. \quad (6)$$

The SGD updating of each dimension of $\boldsymbol{\mu}_{\text{new}}$ is

$$\mu_{\text{new},i} \leftarrow \mu_{\text{new},i} + \eta \cdot \left((y - \sigma(\boldsymbol{\mu}_{\text{new}}^T \mathbf{x})) x_i - q_{\text{old},i} (\mu_{\text{new},i} - \mu_{\text{old},i}) \right) \quad (7)$$

and \mathbf{S}_{new} is given by the inverse of the matrix of second derivatives of the negative log likelihood, which satisfies

$$\mathbf{S}_{\text{new}}^{-1} = -\nabla \nabla \ln p(\boldsymbol{\mu}|\mathbf{x}, y) = \mathbf{S}_{\text{old}}^{-1} + \sigma(\boldsymbol{\mu}^T \mathbf{x})(1 - \sigma(\boldsymbol{\mu}^T \mathbf{x})) \mathbf{x} \mathbf{x}^T.$$




Figure 1: An illustration of the proposed CTR distribution against its parameters μ and q in Eq. (12).

As we follow [40, 13] to assume each w_i as independent with each other and thus \mathbf{S} is diagonal, we have the updating of precision parameters as

$$q_{\text{new},i} = \mathbf{S}_{\text{new},i,i}^{-1} = q_{\text{old},i} + \sigma(\boldsymbol{\mu}^T \mathbf{x})(1 - \sigma(\boldsymbol{\mu}^T \mathbf{x})) x_i^2. \quad (8)$$

3.2 Predicted CTR Distribution

Equipped with a Bayesian logistic regression with a Laplace approximation of the parameter posterior, we are ready to propose a simple yet novel solution for modeling the CTR distribution. Note that the CTR itself is a probability estimation that a click event occurs from an impression with feature \mathbf{x} takes the form

$$\hat{y} = P(y = 1|\mathbf{x}) = \sigma(\mathbf{w}^T \mathbf{x}), \quad w_i \sim N(\mu_i, q_i^{-1}), \quad (9)$$

where \hat{y} denotes the CTR random variable that generates the binary observation y , so our goal is to estimate the distribution of \hat{y} . As introduced in Section 3.1, we assume each w_i is from Gaussian i.i.d., the distribution of $\sum_i w_i x_i$ also follows $N(\sum_i \mu_i x_i, \sum_i q_i^{-1} x_i)$. Thus, we have

$$\hat{y} = \sigma(a), \quad a \sim N\left(\sum_i \mu_i x_i, \sum_i q_i^{-1} x_i\right). \quad (10)$$

Consider that if random variables x and y satisfy $y = g(x)$ and g^{-1} is monotonic and differentiable, we have [36]

$$p_{\hat{y}}(y) = p_x(g^{-1}(y)) \left| \frac{\partial g^{-1}(y)}{\partial y} \right|. \quad (11)$$

In our case, $\sigma^{-1}(\hat{y}) = \ln \hat{y} - \ln(1 - \hat{y})$ is monotonic and differentiable within $(0, 1)$, so we obtain the closed-form of \hat{y} 's p.d.f.

$$p_{\hat{y}}(\hat{y}) = \frac{1}{(\hat{y} - \hat{y}^2) \sqrt{2\pi \sum_i q_i^{-1} x_i}} e^{-\frac{(\sigma^{-1}(\hat{y}) - \sum_i \mu_i x_i)^2}{2 \sum_i q_i^{-1} x_i}}, \quad (12)$$

which provides an explicit CTR p.d.f. To our best knowledge, the above proposed solution has not been studied in previous literature [40, 13, 35]. To understand the nature of the model, Figure 1 plots the CTR distribution against its parameters μ and q , where $\mathbf{x} = \mathbf{1}_{10}$. As observed, the p.d.f. presents a single-peak shape in $[0, 1]$ which is similar with the shapes of beta distributions.

It is straightforward to see from Figure 1 that μ and q jointly determine the peak location and sharpness of the CTR p.d.f. and specifically μ influences more on the peak location while q influences more on the distribution sharpness. To understand how it models the CTR prediction uncertainty/confidence, from Eq. (8), we see each time a data instance with feature \mathbf{x}_i is observed, the precision q_i will be updated with a higher value, which in turn contributes

a sharper CTR p.d.f. in Eq. (12). Therefore, for the ad impression with frequent (similar) features, the predicted CTR is of low uncertainty, and vice versa.

4. RISK-AWARE BIDDING STRATEGIES

With our CTR distribution model in Eq. (12), we next investigate the conditional distribution of a utility given a specific bid price for an input bid request. By considering a risk-aware utility as an optimization target, we are ready to derive the corresponding risk-aware bidding strategy. Note that alternative CTR distribution models can also be incorporated in our solution framework.

Specifically, we start from the theoretic derivation of a Bayesian truth-telling bidding strategy and provide an analysis of its risk in Section 4.1. Then we will propose two solutions, discussed in Sections 4.2 and 4.3 respectively.

4.1 Analysis: Bayesian Truth-telling Bidding

The utility r of an ad impression could be defined based on the advertiser's value v on a specific user action, e.g., click or conversion. For example, if the value v is on each click, given an impression with CTR \hat{y} distribution as in Eq. (12), the utility r and its p.d.f. of this impression are

$$r = v \cdot \hat{y}, \quad p_r(r) = p_{\hat{y}}(\hat{y})/v. \quad (13)$$

Moreover, the cost to win the ad impression comes from the highest bid from other competitors, defined as market price z [5]. The profit of winning this impression is then calculated as the utility r minus the cost z , which is set as the optimization target of performance-driven campaigns [38].

In a general setting, considering both the estimated CTR \hat{y} and cost z are stochastic variables: $\hat{y} \sim p_{\hat{y}}(\hat{y})$ and $z \sim p_z(z)$, the bid optimization problem is to find the optimal bid price to the auction.

Let us first consider a simple case without considering the uncertainty of our estimation, where the goal is to maximize the expected profit $R(b)$ by marginalizing out \hat{y} and z :

$$b^* = \arg \max_b \mathbb{E}[R(b)] \quad (14)$$

$$= \arg \max_b \int_{\hat{y}} \int_{z=0}^b (v \cdot \hat{y} - z) p_z(z) dz \cdot p_{\hat{y}}(\hat{y}) d\hat{y}. \quad (15)$$

Taking the derivative w.r.t. b and set it to 0:

$$\frac{\partial \mathbb{E}[R(b)]}{\partial b} = \frac{\partial}{\partial b} \int_{z=0}^b p_z(z) \int_{\hat{y}} (v \cdot \hat{y} - z) p_{\hat{y}}(\hat{y}) d\hat{y} dz \quad (16)$$

$$= p_z(b) \int_{\hat{y}} (v \cdot \hat{y} - b) p_{\hat{y}}(\hat{y}) d\hat{y} = 0 \quad (17)$$

$$\Rightarrow b^* = \int_{\hat{y}} v \cdot \hat{y} \cdot p_{\hat{y}}(\hat{y}) d\hat{y} = v \cdot \mathbb{E}[\hat{y}] = \mathbb{E}[r], \quad (18)$$

where we see that the optimal bid price is the product of action value and the estimated CTR \hat{y} , which is independent of the market price distribution. If we assume \hat{y} is known and fixed, i.e., $p_{\hat{y}}(\hat{y})$ focuses its mass on a single point, then the optimal bid price is $v \cdot \hat{y}$. Note that the optimality of truth-telling bidding is for one-shot auction. When considering campaign budget and auction volume, a coefficient ϕ is commonly added to Eq. (18) as $\phi \cdot \mathbb{E}[r]$.

Discussion of Risk. The above classic bidding solution is built on maximizing the expectation of the profit $R(b)$, regardless of its uncertainty. However, a potential problem is that there is a chance a bid is won, but $v \cdot \hat{y}$ is less than z , in which case $R(b)$ will be negative. We can obtain the

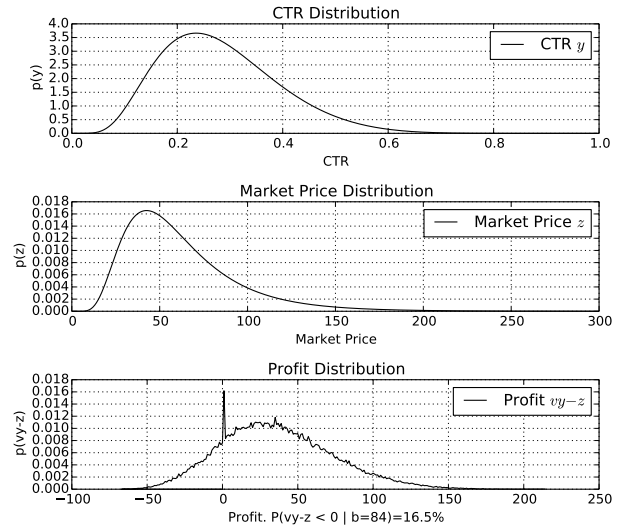


Figure 2: An example of CTR, market price and profit distribution when bidding the expected utility value. The profit p.d.f. on 0 is the probability of losing the auction, resulting in a peak.

probability of such negative profit $P(R(b) < 0)$ as

$$\begin{aligned} P(R(b) < 0) &= \int_0^1 p_{\hat{y}}(\hat{y}) P(b > z > v \cdot \hat{y}) d\hat{y} \\ &= \int_0^1 p_{\hat{y}}(\hat{y}) \int_{v \cdot \hat{y}}^b p_z(z) dz d\hat{y}, \end{aligned} \quad (19)$$

which shows it is of risk to get negative profit whatever the positive bid price. We illustrate our point in Figure 2 using an example. We set the CTR p.d.f. as Eq. (12) with $\sum_i \mu_i x_i = -1$, $\sum_i q_i^{-1} x_i = \frac{1}{3}$; the market price distribution is assumed to be a log-normal p.d.f. with $\mu = 4$, $\sigma = 0.5$; the value per click is set as 300. The expected value of CTR is 0.283, and thus the truth-telling bidding in Eq. (18) would give 84. From our simulation with 10,000 samples, we plot the profit distribution given the truth-telling bid 84 with probability 16.7% the profit is negative. From Eq. (19) we see that the down-side risk, the probability of the CTR lower than the mean \hat{y} , contributes more to the probability of negative profit; the higher standard deviation (or variance) is of \hat{y} , the higher probability of such negative profit may occur in the truth-telling bidding case. Thus, such risk of negative profit should be carefully considered and incorporated into the bidding strategies to help campaigns achieve satisfactory performance with controlled risk.

4.2 Bidding with Value-at-Risk Utility

The utility p.d.f. $p_r(r)$ is of high importance for risk modeling. From Figure 2 we know that due to the uncertainty of utility, bidding with the utility expectation $\mathbb{E}[r]$ will introduce the risk of negative profit, which might be worse than no bid. Let us examine the down-side risk.

LEMMA 1 (CANTELLI'S INEQUALITY). *For a random variable r with mean μ and standard deviation σ , the following inequalities hold:*

$$P(r < \mu - \alpha\sigma) < \frac{1}{1 + \alpha^2} \quad (20)$$

With Lemma 1 we can find a variant of *value at risk* (VaR) [19, 27] for utility

$$\tilde{r} = \mathbb{E}[r] - \alpha \text{Std}[r], \quad (21)$$

which has a guarantee that the real utility is lower than \tilde{r} with a probability smaller than $1/(1 + \alpha^2)$.

Using VaR \tilde{r} as a new risk-aware utility, we set the truth-telling bid at \tilde{r} . Taking Eq. (13) into Eq. (21) gives

$$b_{\text{VaR}}(r) = \tilde{r} = v \cdot (\mathbb{E}[\hat{y}] - \alpha \text{Std}[\hat{y}]), \quad (22)$$

where the p.d.f. of \hat{y} is given in Eq. (12).

When $\alpha > 0$, the bidding strategy is risk-averse, inversely when $\alpha < 0$ it is risk-seeking; the traditional truth-telling bidding (Eq. (18)) is now a special case of $b_{\text{VaR}}(r)$ with $\alpha = 0$, called risk-neutral.

Note that we derive VaR strategy based on Lemma 1 rather than the distribution of \hat{y} (Eq. (12)) directly for the following reasons: (i) Lemma 1 is a general one that can be applied to other CTR distribution models, especially those without analytical forms; (ii) α can be simply set over all bid requests to make the computation more efficient. The uniform α can also be regarded as a risk control parameter.

Considering campaign budget and auction volume, a coefficient ϕ is added to Eq. (22), which is similar to truth-telling bidding. This also applies to the strategy proposed in Section 4.3.

Due to page limit, the detailed realization method and efficiency analysis are provided in the supplementary².

4.3 Bidding for Risk Management of Profit

Another way to approaching risk-aware bidding strategies is to go back to the analysis of campaign profit $R(b)$, which is one of the key performance indicators (KPIs) in RTB.

Mathematically,

$$R(b) = \begin{cases} 0 & b \leq z \text{ (lose)} \\ v \cdot \hat{y} - z & b > z \text{ (win)} \end{cases}. \quad (23)$$

As both CTR \hat{y} and market price z are modeled as stochastic random variables with p.d.f. $p_{\hat{y}}(\hat{y})$ and $p_z(z)$, $R(b)$ can be naturally regarded as a dependent random variable. Again, using Lemma 1, the value-at-risk of profit

$$\tilde{R}(b) = \mathbb{E}[R(b)] - \alpha \text{Std}[R(b)], \quad (24)$$

which has a guarantee that the real profit is lower than $\tilde{R}(b)$ with a probability smaller than $1/(1 + \alpha^2)$.

In this setting, we propose the second risk-aware bidding strategy that generates the bid which yields the maximum value-at-risk of profit:

$$b_{\text{RMP}}(R(b)) = \arg \max_b \mathbb{E}[R(b)] - \alpha \text{Std}[R(b)]. \quad (25)$$

Through analyzing the properties of $R(b)$, we find that $\mathbb{E}[R(b)]$ and $\text{Std}[R(b)]$ have a non-trivial trade-off relation, which can be balanced by b . We solve the optimal b by importing the concept of efficient frontier from finance [16]. Detailed analysis can be found in the supplementary.

5. OFFLINE EMPIRICAL STUDY

In this section, we empirically evaluate our proposed risk-aware bidding strategies on real-world data³.

Before presenting experimental results on real-world data, we firstly studied the effectiveness of the risk-aware bidding strategy on synthetic data. The results indicated that risk-aware bidding strategy does perform better than traditional bidding strategy in some settings. The details is provided in the supplementary.

²Supplementary link: <https://goo.gl/QHUib2>

³Our experiment is repeatable and the code is given below: <https://goo.gl/Wi0lz1>

5.1 Datasets

Two ad log datasets in our empirical study are:

iPinYou is the largest independent programmatic media buying platform in China. The dataset⁴ was released for research after its global competition on RTB algorithms in 2013 [18]. It contains 19.50M ad impressions, 14.89K clicks from 9 campaigns during 10 days in 2013, which involve 16K Chinese Yuan (CNY) expense in total. According to the data publishers, the last three day data of each campaign is split into test data while the former part into training data. The overall effective cost per click (eCPC) is 1.07 CNY on the training data and 1.13 CNY on the test data, which means the overall market competition did not change dramatically during that 10 days.

ANONYM is another mainstream demand-side platform company. This proprietary dataset is mainly used for training the risk-aware models deployed for online A/B testing. Details will be given in Section 6.

Both datasets are with record-per-line format. After ad log joining, each record is formalized as a triple (\mathbf{x}, y, z) , where \mathbf{x} is the high dimensional feature vector for each bid request with the corresponding ad information, y is the user feedback on the ad impression, e.g., the binary click or conversion actions, z is the market price for that auction, i.e., the lowest price to bid in order to win the auction.

5.2 Experiment Protocol

We followed [38] to set up our evaluation procedure. For each campaign bidding strategy with a predefined budget, the optimal parameters $(\boldsymbol{\mu}, \mathbf{S})$ were learned on training data and the hyperparameters (α and ϕ , discussed later) were tuned on validation data, which was the early half split from the test data. Then we replayed the historic bid records to test the performance on the other half of test data.

For each campaign, there was a defined value v for the user action, i.e., a click in our experiment. Following [38] the click value is defined by a proportion of eCPC on training data to imitate the true value set by the advertisers. In this paper, we set the proportion to 100%. Every time the tested bidding agent received a bid request from the ad exchange, it generated a bid with the tested bidding strategy. If the bid price was higher than the historic market price, the agent won the ad impression and paid the market price and then the historic recorded binary user feedback, i.e., click or not, is observed. If there was a user click, the agent made revenue of the click value. The test ended either when there was no more test bid request or when the campaign budget was exhausted, if applicable.

Note that the offline experiment cannot fully simulate the market competition because there is no observed user feedback and market price for historically lost auctions. Nevertheless, our evaluation protocol keeps the bid requests, displayed ads, and auction environment unchanged. We try to answer that under the same context if the campaign were given a different bidding strategy, whether they would be able to get more clicks with the budget limitation.

5.3 Compared Bidding Strategies

The compared bidding strategies are as follows.

- **LR** - Linear Revenue Bidding: the baseline, which is the most widely used bidding function

$$b_{\text{LR}} = \phi \cdot v \cdot \hat{y}, \quad (26)$$

where \hat{y} is the predicted CTR with logistic regression as in Eq. (1). ϕ is the scaling parameter tuned based

⁴Dataset link: <http://data.computational-advertising.org>




Figure 3: Profit v.s. cost of the VaR bidding.

on the market competitiveness [25]. The overall AUC of this LR estimator on iPinYou dataset is 69%.

- **VaR** - Value at Risk Bidding: our first risk-aware bidding strategy proposed in Section 4.2, which bids the value at risk, shown in Eq. (22), where α is the hyperparameter to be tuned on validation data.
- **RMP** - Risk Management of Profit based Bidding: the second risk-aware bidding strategy proposed in Section 4.3, which seeks the optimal bid to maximize the expected profit with its standard deviation as risk constraint, shown in Eq. (25), where α is the hyperparameter to be tuned on validation data.

The two risk-aware bidding strategies rely on the same CTR distribution model proposed in Section 3. The initial prior distribution of \mathbf{w} was given that (i) μ_0 was set as the point estimation generated by logistic regression; (ii) q_0 was set as constant 1. Our model achieves a comparable AUC with standard LR model. Note that for the budgeted bid optimization tasks in Section 5.7, there would also be a hyperparameter ϕ multiplied in VaR (22) and RMP (25) to be tuned to avoid budget under- or over-delivery.

5.4 Evaluation Measures

For a bid optimization task, the major evaluation measures are campaign profit, i.e., the earned total click value minus the cost, and ROI, i.e., the ratio between the profit and the cost. In addition, we also monitor other key metrics, such as cost per thousand impressions (CPM), auction winning rate, CTR and eCPC to gain more insights.

In order to investigate how bidding strategies balance return against risk, we also propose an additional metric: profit - λ cost, which is named as Cost-Penalized Profit (CP-Profit) in this paper. Intuitively, advertisers want to maximize the profit of their performance-driven campaign given the budget, or minimize the advertising cost given the profit, either of which forms a profit/cost tradeoff. Note that this metric was mainly used for model selection, i.e., selecting α and ϕ . As we will show in Section 5.5 about the profit/cost analysis, CP-Profit is an effective metric to balance the profit and the cost to select an optimal model.

5.5 Profit and Cost Analysis

We analyzed the profit and cost of various bidding strategies with different parameter settings. This would help us understand the properties of the bidding strategies and help us select optimal model hyperparameters.

Non-Budgeted Settings. Without budget constraint, the basic LR bidding strategy is truth-telling, i.e. $\phi = 1$. For the risk-aware bidding strategies we varied the value of α . Figure 3 shows the profit and cost change when setting different α 's in VaR Eq. (22), with $\alpha > 0$ as risk-averse, $\alpha < 0$ as risk-seeking and $\alpha = 0$ as risk-neutral. We clearly observed that for all tested campaigns the results generally formed a single-peak concave shape, showing strong trade-off between the profit and cost. Risk-averse strategies intended to yield lower cost than risk-seeking ones, as




Figure 4: Profit v.s. cost of the RMP bidding.




Figure 5: Selected model performance comparison with different λ 's (validation dataset), iPinYou campaign 3427 for RMP strategy with 1/4 budget setting. The models maximizing CP-Profit with $\lambda = 0, 0.2, 0.4$ are selected, which is highlighted with red circles. Note that LR and risk-aware models are selected separately and then plotted together.

the risk-averse strategies would always bid lower than the risk-neutral one, and the latter further bid lower than the risk-seeking ones. The performance-driven campaigns pursue high profit and low cost. Thus if a bidding strategy yields higher profit and lower cost than another one, we say the former strategy *dominates* the latter one. From the results we saw that almost all the risk-seeking strategies were dominated by part of the risk-averse strategies. That is to say, for almost every risk-seeking setting with $\alpha < 0$, there always exists one or more risk-averse setting with $\alpha > 0$ that yields equal or higher profit with lower cost.

Figure 4 illustrates the profit and cost trade-off yielded by RMP Eq. (25) with different risk settings. We also observed that, similarly, the risk-averse strategies tended to yield lower cost than the risk-seeking ones. The risk-neutral strategy acted as a splitting point between the risk-averse and risk-seeking ones.

Budgeted Settings. In practice, advertisers tend to set up campaign budget constraints. With the budget constraint, the basic LR bidding strategy is not necessarily truth-telling, i.e., $\phi \neq 1$ in Eq. (26) [25]. We performed an analysis on the profit and cost trade-off for different budget constraints. For each campaign, we followed [38] to use $1/2, 1/4, \dots, 1/32$ of the original total cost in the test data as the budget. For LR we varied the value of ϕ while for VaR and RMP we varied both ϕ and α to obtain different profits and costs. As an example, the results of iPinYou campaign 3427 for RMP with 1/4 budget setting are shown in Figure 5. We saw that part of the risk-averse strategies dominated other strategies, which was in accordance with the ones without budget constraints.

Distinction with Lowering Bidding. The performance of risk-averse bidding strategies was not due to lowering bidding. As we know, for a truth-telling linear bidding, if we lower the bid price uniformly, it would result in cost decreasing, profit decreasing and ROI increasing. Unlike




Figure 6: Model selection. Colored dashed lines stand for $\lambda = 0$ (blue), 0.2 (green), 0.4 (red) respectively.




Figure 7: Selected model performance comparison with different λ 's (validation dataset).

this, our risk-averse bidding strategies lowered the bid price of uncertain requests and raised the bid price of confident requests, which led to higher profit and higher ROI with invariant cost. It can be verified in Figure 5. For example, the left circled LR model lowered the bid price of the right circled LR model by applying a smaller ϕ , so it achieved lower cost, lower profit and higher ROI. In the meantime, the risk-averse models above the right circled LR model, which had a positive α and a larger ϕ , achieved the same cost, higher profit and higher ROI.

5.6 Bidding without Budget Constraints

In order to select different risk-aware models (i.e. the strategies' hyperparameter α) with various risk-return balanced metrics, we used the proposed CP-Profit metric. Specifically, with a fixed value of λ , the contour of a specific CP-Profit is a straight line in the cost-profit plot. The highest CP-Profit is achieved by the tangent point of the contour and the cost-profit points. Figure 6 provides example tangent points with different λ 's on the cost-profit plot. For a fixed λ , we then obtained an optimal α that maximizes the CP-Profit.

Figure 7 plots the performance of the tested bidding strategies against CP-Profit metrics with various λ in validation data. We found that: (i) The risk-aware strategies VaR and RMP outperformed the baseline LR on all studied campaigns with all λ settings. (ii) Specifically, when $\lambda = 0$, i.e., the evaluation metric was purely the profit, VaR and RMP outperformed LR with 50.7% and 25.1% improvements. Also note that there were 89% selected models with positive α 's, which suggested when there is a fixed volume of bid requests and no limited budget, the risk control helps advertisers make higher profit by spending less money on opportunities with high risk. (iii) For metrics with $\lambda > 0$, the performance gained of VaR and RMP over LR were larger. As we observed that the CP-Profit curves of VaR and RMP were convex while the ones of LR was a straight line, which meant the risk-aware strategies provided higher ROI when the model selection metrics became more conservative, i.e., with higher λ .

With the model selection on validation data, we compared the strategies on test data. The overall performance across 9 iPinYou campaigns is shown in Figure 8. We ob-




Figure 8: Overall non-budgeted test performance.

served that: (i) Most risk-aware bidding strategies yielded significantly higher profit than the baseline LR, which demonstrated the effectiveness of incorporating risk modeling in bid optimization. (ii) All VaR strategies yielded significantly higher ROI compared to LR, which meant bidding at risk helps advertisers avoid from wasting money on risky opportunities. It turns out that such risky opportunities brought little revenue so avoiding them led to a lower cost and higher profit, and thus much higher ROI. (iii) RMP strategies stably yielded about 10% profit more than LR while the ROI was lower than LR. This is because RMP seeks the bid price which directly maximizes the value at risk of profit, which resulted in consistent profit gains. On the other hand, RMP does not explicitly control the bid price like VaR, thus it did not yield lower CPM or cost. As a result, its ROI was not higher than LR. (iv) All the risk-aware bidding strategies brought higher CTR than the baseline LR, which might be counterintuitive. Actually, the proposed strategies allocate the budget from the high-uncertainty cases to low-uncertainty ones, which does not mean the average CTR should get lower.

5.7 Bidding with Budget Constraints

To follow the previous section, we sought the tangent points which maximized CP-Profit with different λ . The selected models are highlighted with red circles in Figure 5. There might be overlapped red circles as the CP-Profit with different λ 's still selected the same model.

Figure 9 provides the overall performance over 9 iPinYou campaigns with budget constraint (1/2 budget setting). The baseline is LR with $\lambda = 0$, i.e., simply selecting ϕ yielding the highest profit on validation dataset. We found that: (i) For small λ settings, VaR and RMP were both better than LR on profit, which verifies the claim that the risk-aware bidding strategies successfully yield higher profit via controlling the risk. Specifically, risk-aware bidding strategies performed better on yielding profit for the campaigns with poorer LR estimators. (ii) All proposed methods, i.e., LR with non-zero λ model selection and VaR/RMP with various λ 's, yielded higher ROI than the baseline LR with $\lambda = 0$. Some strategies (VaR and RMP with $\lambda = 0, 0.2$) yielded both higher profit and higher ROI. (iii) All risk-aware bidding strategies provided higher CTR (except for the conservative VaR with $\lambda = 0.4$), which meant the proposed algorithms filtered out the low-value cases, which were always with high uncertainty. (iv) For CPM and winning rate, VaR reduced the CPM by lowering bids on uncon-




Figure 9: Performance with budget constraint (1/2 budget).




Figure 10: Online results from ANONYM DSP.

fidant cases and highering the bids on confident ones (via tuning ϕ); RMP did not guarantee any low bid as it just sought the bid yielding the highest value of risk-controlled profit.

6. ONLINE DEPLOYMENT AND A/B TEST

Online Environment. We have deployed our bidding strategies on ANONYM Platform, which is a mainstream DSP in the global RTB market. The total daily bid request volume received by ANONYM was about 5 billion. Among the received bid requests, about 200M (4%) met ANONYM’s bid target rules of cost-per-click (CPC) campaigns, for which ANONYM would return a valid bid. Each of our tested strategies was allocated with 1.5% bid request volume from 10 CPC campaigns, i.e. about 3M bid opportunities daily. For training our models, we collected the impression and click data of 7 consecutive days just before the A/B testing in Jan. 2016, which contained 424M ad impressions, 532K clicks and 47.2K US dollars (USD) expense. The eCPC on the training data is 0.087 USD. A 2% negative instance down sampling was performed by ANONYM system. The advertisers of tested CPC campaigns paid a fixed amount for every user landing action, i.e., valid click. The landing value v was set to be 0.054 USD, which was the average CPC of all campaigns.

Test Setting. We tested two strategies for 7 days in Jan. 2016: (i) The baseline LR strategy. It consists of a logistic regression whose AUC is 88% and a linear bidding function with fixed $\phi = 0.56$ set by ANONYM operation staff. (ii) The proposed VaR strategy. The parameter α acquired

from the training was a positive value that slightly varied daily, which means the trained bidding strategy risk-averse. The parameter ϕ was set to be greater than LR. Thus the two parameters had opposite effects on bid price, which resulted in an average bid price close to that of LR. Strategies were tested with the same bid volume constraint. The bid volume for each strategy was 19.5M. The whole test live volume contained 39M auctions resulting in 3.3M impressions and 7.4K user landings. VaR spent 12% more budget than LR since ϕ was learned from training thus there was no guarantee of equivalent budget spending for A/B test.

Result Discussion. The online results are presented in Figure 10. We show the performance over six metrics. From the results, we found that: (i) Our VaR strategy outperformed baseline LR on ROI by about 5%, which demonstrated the efficacy of our risk-aware bidding strategy. (ii) VaR achieved 17.5% higher profit than LR. Although this was partly due to the more budget spent by VaR, the more important fact was that VaR strategy could help advertiser spend such more budget at a equivalent or even higher ROI within the same bid request volume. Note that advertisers determine their budget usually depending on ROI and ROI is always negative correlative with budget, thereby if a strategy allows an additional amount of budget spent on the same ROI, advertisers would like to increase their budget by that amount. (iii) VaR won much fewer auctions than LR. According to our setting, VaR bid higher for confident cases and lower for risky cases. Its low winning rate result indicated that the extra confident cases VaR gained were fewer than the risky cases VaR lost. (iv) As for CTR, VaR achieved 64.4% higher CTR than baseline LR, which demonstrated the high efficiency of risk-aware strategies in finding high quality ad display opportunities. (v) VaR got higher CPM than LR. This indicated that our strategy tended to target at impressions with higher quality at the cost of higher price. Although VaR strategy achieved higher CPM, it also achieved higher CTR and the co-effect was reflected in low eCPC and high ROI, which was beneficial.

7. CONCLUSION

In this paper, we presented a solution for modeling the uncertainty of CTR estimation in RTB display advertising. With a Bayesian logistic regression CTR estimator, we obtained a closed form of CTR distribution density. On the basis of the distributions of CTR and the market price, two risk-aware bidding strategies are formulated: the first one (VaR) bids the value at risk of the estimated utilities, while the second one (RMP) seeks the optimal bids that maximizes a lower bound of profit with a controlled risk. Our risk-return analysis and offline experiment demonstrated 15.4% profit gain over a linear optimized bidding strategy. To test the applicability of our risk-aware bidding strategies in a live setting, the bidding strategies have been deployed on an operational platform. 17.5% profit gain over the baseline was observed in a 7-day online A/B test.

For future work, we will analyze the market equilibrium if more advertisers adopt the risk-aware bidding strategies. We also plan to further explore applications from the proposed CTR distribution model and the risk-aware bid optimization framework. For an ad placement with multi-frame dynamic creatives, a portfolio optimization [11] can be performed to select the optimal creatives combination to yield the highest profit with a controlled risk. Moreover, our CTR distribution model is potential to be dynamically updated to evaluate users’ interest on an advertised product after repeated ad displays over time.

8. REFERENCES

- [1] Daily NYSE group volume. <http://goo.gl/2EflkC>. Accessed: 2016-02.
- [2] Fikisu mobile DSP. <https://goo.gl/M39SjZ>. Accessed: 2016-02.
- [3] iPinYou statistics. <http://goo.gl/Jd4xIV>. Accessed: 2016-02.
- [4] The Shanghai Stock Exchange data. <https://goo.gl/F12JxZ>. Accessed: 2016-02.
- [5] K. Amin, M. Kearns, P. Key, and A. Schwaighofer. Budget optimization for sponsored search: Censored learning in MDPs. *UAI*, 2012.
- [6] M. Avellaneda and S. Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008.
- [7] C. M. Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [8] Y. Chen, P. Berkhin, B. Anderson, and N. R. Devanur. Real-time bidding algorithms for performance-based display ad allocation. *KDD*, 2011.
- [9] M. Crouhy, D. Galai, and R. Mark. *The essentials of risk management*, volume 1. McGraw-Hill New York, 2006.
- [10] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *NBER*, 2005.
- [11] E. J. Elton, M. J. Gruber, S. J. Brown, and W. N. Goetzmann. *Modern portfolio theory and investment analysis*. JWS, 2009.
- [12] A. Ghosh, B. I. Rubinstein, S. Vassilvitskii, and M. Zinkevich. Adaptive bidding for display advertising. In *WWW*, 2009.
- [13] T. Graepel, J. Q. Candela, T. Borchert, and R. Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. In *ICML*, pages 13–20, 2010.
- [14] D. J. Hand and K. Yu. Idiot’s bayes not so stupid after all? *International statistical review*, 69(3):385–398, 2001.
- [15] X. He, J. Pan, O. Jin, T. Xu, B. Liu, T. Xu, Y. Shi, A. Atallah, R. Herbrich, S. Bowers, et al. Practical lessons from predicting clicks on ads at facebook. In *ADKDD*, 2014.
- [16] J. Hull. *Risk Management and Financial Institutions*, volume 733. Wiley, 2012.
- [17] K.-C. Lee, B. Orten, A. Dasdan, and W. Li. Estimating conversion rate in display advertising from past performance data. In *KDD*, 2012.
- [18] H. Liao, L. Peng, Z. Liu, and X. Shen. ipinyou global rtb bidding algorithm competition dataset. In *ADKDD*, 2014.
- [19] T. J. Linsmeier and N. D. Pearson. Value at risk. *Financial Analysts Journal*, 56(2):47–67, 2000.
- [20] H. Markowitz. Portfolio selection. *The JF*, 1952.
- [21] H. B. McMahan et al. Ad click prediction: a view from the trenches. In *KDD*, 2013.
- [22] J. Mun. *Modeling risk: Applying Monte Carlo simulation, real options analysis, forecasting, and optimization techniques*, volume 347. John Wiley & Sons, 2006.
- [23] S. Muthukrishnan. Ad exchanges: Research issues. In *WINE*. 2009.
- [24] R. J. Oentaryo, E. P. Lim, D. J. W. Low, D. Lo, and M. Finegold. Predicting response in mobile advertising with hierarchical importance-aware factorization machine. In *WSDM*, 2014.
- [25] C. Perlich, B. Dalessandro, R. Hook, O. Stitelman, T. Raeder, and F. Provost. Bid optimizing and inventory scoring in targeted online advertising. In *KDD*, 2012.
- [26] M. Richardson, E. Dominowska, and R. Ragno. Predicting clicks: estimating the click-through rate for new ads. In *WWW*, pages 521–530. ACM, 2007.
- [27] R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.
- [28] W. F. Sharpe. The sharpe ratio. *Streetwise—the Best of the Journal of Portfolio Management*, 1998.
- [29] J. Shen, B. Orten, S. C. Geyik, D. Liu, S. Shariat, F. Bian, and A. Dasdan. From 0.5 million to 2.5 million: Efficiently scaling up real-time bidding. In *ICDM*, 2015.
- [30] Y. Shi, X. Zhao, J. Wang, M. Larson, and A. Hanjalic. Adaptive diversification of recommendation results via latent factor portfolio. In *SIGIR*, 2012.
- [31] I. Trofimov, A. Kornetova, and V. Topinskiy. Using boosted trees for click-through rate prediction for sponsored search. In *WINE*, page 2. ACM, 2012.
- [32] J. Wang and B. Chen. Selling futures online advertising slots via option contracts. In *WWW*, 2012.
- [33] J. Wang and J. Zhu. Portfolio theory of information retrieval. In *SIGIR*, 2009.
- [34] L. Wang, P. N. Bennett, and K. Collins-Thompson. Robust ranking models via risk-sensitive optimization. In *SIGIR*, 2012.
- [35] X. Wang, W. Li, Y. Cui, R. Zhang, and J. Mao. Click-through rate estimation for rare events in online advertising. *Online Multimedia Advertising: Techniques and Technologies*, 2010.
- [36] L. Wasserman. *All of statistics: a concise course in statistical inference*. Springer Science & Business Media, 2013.
- [37] S. Yuan, J. Wang, and X. Zhao. Real-time bidding for online advertising: measurement and analysis. In *ADKDD*, 2013.
- [38] W. Zhang and J. Wang. Statistical arbitrage mining for display advertising. In *KDD*, 2015.
- [39] W. Zhang, S. Yuan, and J. Wang. Optimal real-time bidding for display advertising. In *KDD*, 2014.
- [40] Y. Zhang, D. Wang, G. Wang, W. Chen, Z. Zhang, B. Hu, and L. Zhang. Learning click models via probit bayesian inference. In *CIKM*, 2010.
- [41] H. Zhu, H. Xiong, Y. Ge, and E. Chen. Mobile app recommendations with security and privacy awareness. In *KDD*, 2014.