



Interfaces with Other Disciplines

Combining guaranteed and spot markets in display advertising: Selling guaranteed page views with stochastic demand

Bowei Chen^{a,*}, Jingmin Huang^b, Yufei Huang^{c,d}, Stefanos Kollias^e, Shigang Yue^e^aAdam Smith Business School, University of Glasgow, UK^bSchool of Computing Science, University of Glasgow, UK^cTrinity Business School, Trinity College Dublin, Ireland^dSchool of Management, Northwestern Polytechnical University, China^eSchool of Computer Science, University of Lincoln, UK

ARTICLE INFO

Article history:

Received 2 February 2018

Accepted 31 July 2019

Available online 5 August 2019

Keywords:

Dynamic programming

Display advertising

Programmatic guaranteed

Real-time bidding

Revenue maximisation

ABSTRACT

While page views are often sold instantly through real-time auctions when users visit websites, they can also be sold in advance via guaranteed contracts. In this paper, we present a dynamic programming model to study how an online publisher should optimally allocate and price page views between guaranteed and spot markets. The problem is challenging because the allocation and pricing of guaranteed contracts affect how advertisers split their purchases between the two markets, and the terminal value of the model is endogenously determined by the updated dual force of supply and demand in auctions. We take the advertisers' purchasing behaviour into consideration, i.e., risk aversion and stochastic demand arrivals, and present a scalable and efficient algorithm for the optimal solution. The model is also empirically validated with a commercial dataset. The experimental results show that selling page views via both channels can increase the publisher's expected total revenue, and the optimal pricing and allocation strategies are robust to different market and advertiser types.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Display advertising is one of the most popular forms of online marketing. It uses the Internet and the World Wide Web as an advertising medium, and when users visit websites, the promotional messages (i.e., the *ads*), appear on the pages. They usually come in terms of rectangular images or photos placed on a web page either above, below or on the sides of the page's main content and are linked to other web pages. Online publishers make profits by selling the page views, namely the *impressions*, through two channels: (i) selling them in advance via contracts; or (ii) auctioning them off in real time when users visit the web pages. The former is called *guaranteed contracts* (or *reservation contracts*) while the latter is called *real-time bidding* (RTB). Over the past decades, RTB has become the widely used sales model for display advertising, in which advertisers come to a common marketplace, i.e., ad exchange, to compete for impressions from their targeted users (Mansour, Muthukrishnan, & Nisan, 2012; Muthukrishnan, 2009). It is real-time, impression-level and auction-based, thus has

achieved a significant level of automation, integration and user-targeting (Sun, Dawande, Janakiraman, & Mookerjee, 2016; Yuan, Wang, Li, & Qin, 2014).

Although RTB is more widely used, guaranteed contracts in fact have a longer history. In 1994, wire.com signed fourteen contracts with companies, such as AT&T, Club Med and Coor'z Zima, being recognised as the start of display advertising (DoubleClick, 2005). A guaranteed contract is an agreement and it is usually negotiated privately between a publisher and an advertiser for bulk sales. Only a small portion of impressions on the market is sold through guaranteed contracts but they bring in much more revenue than RTB (eMarketer, 2013). In order to meet the demand for automation due to the huge number of site visits, standardised guaranteed contracts have been recently discussed. This is known as *programmatic guarantee* (PG). In essence, PG is a sales system that sells future impressions via standardised guaranteed contracts in addition to RTB (OpenX, 2013). Examples include Google DoubleClick's Programmatic Guaranteed, AOL's Programmatic Upfront and Rubicon Project's Reserved Premium Media Buys. Recent studies have investigated PG from different perspectives and we provide an extensive review in Section 2. These studies aim to answer the following two main questions: (i) how many future impressions should be allocated to guaranteed contracts? and (ii) how to price the guaranteed contracts?

* Corresponding author.

E-mail addresses: bowei.chen@glasgow.ac.uk (B. Chen), jingmin.huang@glasgow.ac.uk (J. Huang), yufei.huang@tcd.ie (Y. Huang), skollias@lincoln.ac.uk (S. Kollias), syue@lincoln.ac.uk (S. Yue).

In this paper, we use a revenue maximisation model to study how to sell impressions using RTB and PG. More specifically, some impressions are sold in advance via standardised guaranteed contracts before the delivery day while the rest of the impressions are auctioned off in RTB using the second-price auctions (Ben-Zwi, Henzinger, & Loitzenbauer, 2015). We focus on one ad slot and study how the estimated total impressions in a future period, i.e., the ad delivery day, should be allocated and priced algorithmically between guaranteed and spot markets. For example, AOL sells the impressions from the top banner of its homepage on Christmas Day. AOL needs to decide on how many impressions should be sold a couple of days, weeks or months before Christmas via guaranteed contracts and at what prices. Unlike the traditional way of selling guaranteed contracts, in our model, there is no negotiation process between the publisher and the advertiser. Instead, the guaranteed contract price is posted in a common marketplace. Advertisers can monitor the price trend over time and purchase the needed impressions directly at the corresponding prices prior to the delivery day. Based on auction theory and operations research studies, we also consider the distinct characteristics of advertisers, i.e., risk aversion and stochastic arrivals, and then propose an algorithm to find the optimal allocation and pricing strategy based on the solution framework to the Knapsack problem. Our solution is a greedy algorithm and there are two significant differences to the Knapsack problem solution. First, for a given allocated purchase demand at a time point, the corresponding guaranteed contract price is obtained by considering advertiser's purchase behaviour. Second, the terminal value of RTB is not certain, which is determined by the updated dual force of supply and demand. The proposed model is further examined with an RTB dataset from a UK supply-side platform (SSP).¹ Our results show that introducing guaranteed contacts in addition to RTB increases the publisher's expected total revenue. Specifically, for ad slots with a high competition level, the guaranteed contract price significantly increases over time and the publisher allocates a large percentage of future impressions into guaranteed contracts, consequently, the revenue is mainly collected by PG. For ad slots with a low competition level, the guaranteed contract price increases steadily and less contracts are sold. However, our results show that the publisher's revenue from those ad slots can be significantly increased because there is a greater margin to be optimised. We further examine the robustness of our model by considering two extensions: (i) incorporating uncertainty into demand and supply of the impressions; and (ii) segmenting advertisers based on their valuations. Our analysis shows the model is robust under different conditions and the publisher can always obtain a higher total revenue when selling through both channels.

Our research makes the following contributions. First, different from the existing work which focuses on either guaranteed contracts or RTB, this paper is among the first to introduce a unified framework that combines PG and RTB simultaneously in display advertising. Second, we also study the interaction of pricing and allocation decisions on guaranteed contracts, thus can provide further insights to the existing literature that only focuses on either pricing or allocation. Third, different from the widely used conventional methodologies in marketing and operational studies, this paper proposes a data-driven analytical model. Our solution to the revenue maximisation problem is simple, efficient and scalable. The insights from our model are further validated using a commercial dataset. The robustness of the results indicates the potential usefulness of our model in practice.

The rest of the paper is structured as follows. Section 2 reviews the related literature. Section 3 discusses the model, in-

cluding problem formulation, model assumptions and our solution. Section 4 describes the used dataset and experimental settings. Section 5 presents our experimental results and Section 6 concludes the paper.

2. Related work

Our paper focuses on selling impressions via both RTB and PG, thus it naturally lies in the interface between mechanism design for online advertising auction and revenue management with dynamic pricing.

Mechanism design for online advertising has been extensively studied in the literature. Many discussions have been centred around search advertising auctions, such as the generalised first-price auction (Edelman, Ostrovsky, & Schwarz, 2007), the generalised second-price auction (Edelman et al., 2007; Lahaie & McAfee, 2011; Lahaie & Pennock, 2007; Varian, 2007), the Vickrey-Clarke-Groves (VCG) auction (Parkes, 2007; Varian, 2009; Varian & Harris, 2014) and the optimal auction (Feldman, Mirrokni, Muthukrishnan, & Páii, 2010; Ostrovsky & Schwarz, 2011; Thompson & Leyton-Brown, 2013) which extends Myerson's optimal auction for a single indivisible good (Myerson, 1981). Most of these studies on advertising auctions examine the properties of an auction model with respect to incentive compatibility, expected revenue, individual rationality, and computational complexity. In this paper, we focus on display advertising with the second-price auction. It has a simplified scenario in auction mechanism design because: (i) the measurement model is based on ad display rather than click, so click-through rate is not a major factor; (ii) ad slots on the same web page for a single page view is auctioned off separately so that each independent auction is a single-item auction.

Our paper is also related to the literature in revenue management (Bitran & Caldentey, 2003; Talluri & van Ryzin, 2005), in which many studies focus on how a seller uses dynamic pricing models to produce or offer a menu of products or services to its customers. For example, Gallego and van Ryzin (1994) used intensity control to sell a given stock of products by a deadline when demand is price sensitive and stochastic and the seller's objective is to maximise his expected revenue. Their model fits many applications such as single-route flight tickets selling and hotel rooms booking. Anjos, Cheng, and Currie (2004, 2005) proposed a dynamic pricing framework for selling flight tickets under the assumption of static demand. Our problem setting for PG is similar to the existing literature, however, the terminal value in our case is uncertain because the remaining impressions are auctioned off in RTB.

Existing literature has also studied selling products or services via both auctions and posted prices. For example, Caldentey and Vulcano (2007) discussed a problem of two channels selling, where the products can be sold through either an auction or an alternative channel with a posted price. They considered two scenarios of this dual-channel optimisation problem: in the first scenario, the posted price is an external channel run by another company; in the second scenario, the seller manages both auction and posted price channels. The second scenario is similar to our model setting. However, their discussion is mainly about the static posted price and they assume that the original values are uniformly distributed and there is no penalty cost. Gao, Shou, Chen, and Huang (2016) studied a hybrid model that unifies both future and spot markets for dynamic spectrum access, in which buyers can purchase under-utilised licensed spectrum either through predefined contracts or through spot transactions with a VCG-like auction model. Their work is similar to ours, however, the seller does not optimise the contract price dynamically.

¹ Supply-side platforms (SSPs) are intermediaries who help publishers sell impressions in RTB. Through SSPs, publishers are able to connect with advertisers from various ad exchanges and networks.

There are also a number of studies discussing the dual-channel problem in the context of display advertising. Feldman, Korula, Mirrokni, Muthukrishnan, and Pál (2009) proposed a selection and matching algorithm for display ads, in their paper, the publisher's objective is not only to fulfil guaranteed contracts but also to deliver well-targeted impressions to advertisers. Ghosh, McAfee, Papineni, and Vassilvitskii (2009) proposed that a publisher can act as a bidder to bid for guaranteed contracts – the allocation of impressions becomes the competition between the publisher and other advertisers so that impressions can be allocated to auctions only if advertisers' bids are high enough. Roels and Fridgeirsdottir (2009), Salomatin, Liu, and Yang (2012), Balseiro, Feldman, Mirrokni, and Muthukrishnan (2014) and Chen (2017) discussed the optimal allocation using stochastic control models. Bharadwaj et al. (2012) proposed a lightweight allocation framework which lets the real servers to allocate ads efficiently and with little overhead. Chen (2016) investigated the dynamic reserve prices of guaranteed contracts based on RTB. However, the model does not give the optimal solution. Hojjat, Turner, Cetintas, and Yang (2014) discussed an idea in the allocation and serving of display ads by using predetermined fixed length streams of ads. Their framework introduces user-level perspective into the common aggregate modelling of the ad allocation problem. Zhang et al. (2017) proposed a consumption minimisation model, in which the primary objective is to minimise the user traffic consumption to satisfy all contracts. Cancellations have also been discussed in some studies, namely, the publisher can cancel a guaranteed contract later if he agrees to pay a penalty (Babaioff, Hartline, & Kleinberg, 2009; Constantin, Feldman, Muthukrishnan, & Pál, 2009).

Guaranteed contract pricing has also been discussed in several recent studies. Bharadwaj et al. (2010) presented two algorithms to compute the price of a guaranteed contract based on the statistics of users' visits to the web pages. Najafi-Asadolahi and Fridgeirsdottir (2014) and Fridgeirsdottir and Najafi-Asadolahi (2018) used queueing systems and discussed two different pricing schemes for a publisher who promises to deliver a certain number of clicks or impressions on the ads posted, where uncertain demand, traffic and click behaviour are considered. Wang and Chen (2012), Chen and Wang (2015) and Chen and Kankanhalli (2019) discussed several pricing methods for various flexible guaranteed contracts tailored to display advertising, called ad options. The ideas came from financial and real options (Constantinides & Malliaris, 2001). Simply, if an advertiser pays a small fee to buy an ad option, he is guaranteed a priority buying right but not an obligation of his targeted future impressions. He can then decide to pay the fixed price in the future to advertise.

Our research in this paper concerns both pricing and allocation so the optimal solution includes and reflects their interaction effects. Our problem setup is similar to Chen, Yuan, and Wang (2014). However, both the model and the analysis are significantly different from theirs in three important aspects. First, we consider stochastic demand for buying guaranteed contracts. We use a Poisson process to model the arrival of advertisers and allow unfulfilled demand to be backlogged. While Chen et al. (2014) assumes the demand for advertising in the future period can be shifted in advance by using a deterministic exponential decay function, the unfulfilled demand in their setting is not explicitly considered at later time points. Second, we devise an optimal pricing and allocation solution to maximise the publisher's expected total revenue. Different from Chen et al. (2014) where the optimal solution is linearly searched, our solution is a greedy algorithm, and is relatively scalable and efficient. Third, we further analyse the model's robustness by incorporating supply and demand uncertainty and customising optimal pricing and allocation for different advertiser segments.

3. Model

Fig. 1 presents a schematic view of the model. It demonstrates how a publisher can sell impressions from a specific ad slot of a publisher between guaranteed and spot markets. Specifically, impressions can be sold in advance via standardised guaranteed contracts and the remaining impressions will be auctioned off in RTB when online users visit the corresponding hosting web page. Let $[0, T]$ be the selling period of guaranteed contracts and $[T, \tilde{T}]$ be the period that impressions are created and auctioned off in RTB. We further use t_n , $n = 0, 1, \dots, N$, to denote N equally spaced discrete time points during the selling period. The relationship between the discrete-time and continuous-time notations is: $t_0 = 0$, $t_N = T$, and $t_{\tilde{N}} = \tilde{T}$. Suppose that the total supply of impressions S in the future period $[t_N, t_{\tilde{N}}]$ is well estimated, the publisher needs to decide how many future impressions to sell in advance through guaranteed contracts during $[t_0, t_N]$ and what price should be charged at each specific time point t_n prior to RTB. The publisher's decision making takes the buying behaviour of advertisers into account. We assume that the total demand of future impressions Q can be well estimated but advertisers arrive stochastically over time prior to the delivery day. For simplicity and without loss of generality, we assume that each advertiser has unit demand (i.e., one impression) so each guaranteed contract is a standardised contract only containing a single impression. This setting is reasonable because impressions are auctioned off individually in RTB (Ben-Zwi et al., 2015) and the setting can be easily extended to bulk sale in practice. We assume the total demand exceeds the total supply to ensure there are competitions among advertisers in the future auctions. Otherwise, no guaranteed contracts would be purchased in advance because they can obtain the needed impressions at very low prices or even reserve prices in auctions (Yuan, Wang, Chen, Mason, & Seljan, 2014). For the reader's convenience, a notation table is provided in Appendix A. In our model setting, the publisher's goal is to sell S impressions from a specific delivery period to Q advertisers with unit demand and stochastic arrival to maximise its revenue. By allowing the publisher to sell these impressions via both guaranteed contract before the delivery period and RTB in the delivery period, we investigate two decisions that the publisher needs to make: (i) how to allocate impressions between the two channels? and (ii) how to set the unit price for the impressions sold via standardised guaranteed contracts at different time points before RTB?

3.1. Stochastic demand arrivals and purchase behaviour

A distinguishing feature of our model is that we consider stochastic demand in buying guaranteed contracts. For ease of exposition, we assume advertisers arrive following a homogeneous Poisson process with a constant λ .² Let $\Delta t = t_n - t_{n-1}$ and $f(t_n)$ be the expected arrivals in the period Δt so $f(t_n) = \lambda \Delta t$. Once advertisers arrive, they will last up to time t_n if their demand is not fulfilled, and we normalise the waiting cost to zero.³ The cumulative expected total demand of buying in advance up to time t_N

² Poisson process is widely adopted in the literature mainly because of the memoryless property of the exponential inter-arrival distribution (Aviv & Pazgal, 2008; McGill & van Ryzin, 1999). It captures the randomness between arrivals and helps us formulate the dynamic pricing setting for the guaranteed contracts with RTB. Our model framework can also accommodate other demand arrival processes, such as non-homogeneous Poisson process. This may make the model expression more complicated but the main insights will still hold, and we believe this is a relatively minor technical concern.

³ In our model, we normalise the waiting cost to zero, because the impressions are all delivered in the same time, i.e., in period $[t_N, t_{\tilde{N}}]$. We do not consider the case where the advertisers may strategically delay their purchases to wait for a lower price in future time periods (prior to the delivery), which will require stronger assumptions regarding how the advertisers hold beliefs on future prices

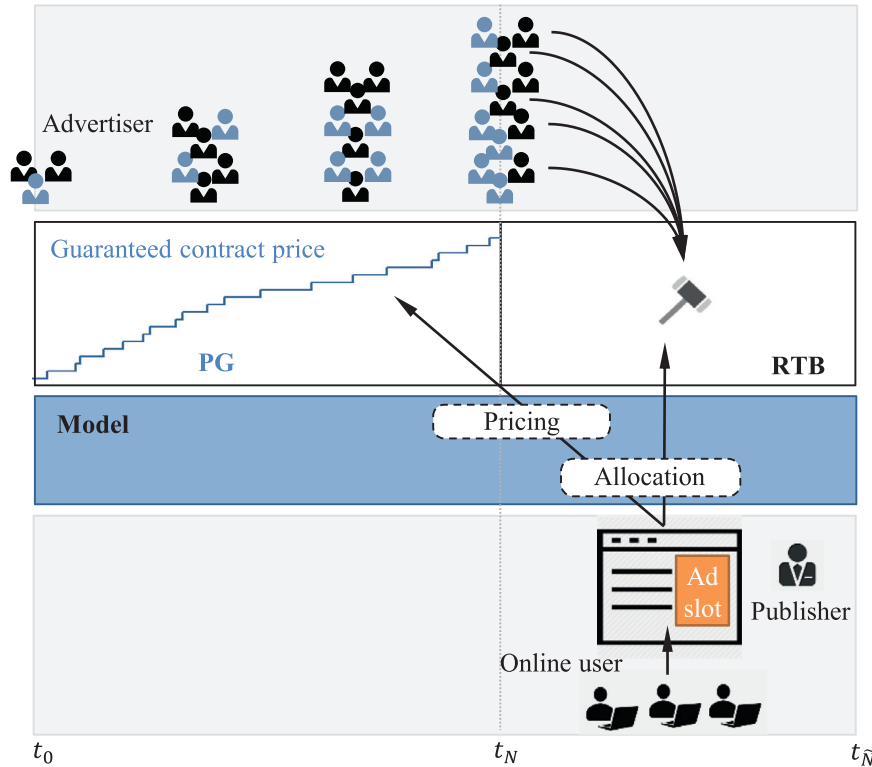


Fig. 1. Schematic view of the proposed model for display advertising: $[t_0, t_N]$ is the time period to sell the guaranteed contracts containing the impressions which will be created in the future period $[t_N, t_{\bar{N}}]$; advertisers' demand of advertising in $[t_N, t_{\bar{N}}]$ arrives sequentially over time in $[t_0, t_N]$; and the unfulfilled demand will join RTB in $[t_N, t_{\bar{N}}]$.

should not be larger than the estimated total demand in the delivery period, that is, $\lambda \leq Q / (\sum \Delta t) = Q / T$.

As the auction outcome is uncertain, we assume advertisers are risk-aware (Chen et al., 2014; Radovanovic & Heavlin, 2012) and we follow the framework proposed by Anjos et al. (2004, 2005) to assume that a certain percentage of arrived advertisers would like to buy guaranteed contracts. Time and price are two key factors of an advertiser's buying decision on the guaranteed contract. Therefore, a ratio function $\theta(t, p(t))$ is used to represent the proportion of those who want to buy an impression in advance at time t and at price $p(t)$, satisfying the following properties:

$$\theta(t, p) \geq \theta(t, p^*), \quad \text{for } 0 \leq p \leq p^*, 0 \leq t \leq T, \tag{1}$$

$$\theta(t, p) \geq \theta(\tau, p), \quad \text{for } 0 \leq \tau \leq t \leq T, 0 \leq p, \tag{2}$$

$$\theta(t, 0) = 1, \quad \text{for } 0 \leq t \leq T. \tag{3}$$

Eq. (1) shows that at the same time point, more advertisers are willing to buy a guaranteed contract when the price is lower; Eq. (2) indicates that more advertisers are willing to buy a guaranteed contract when it is closer to the end of time horizon; Eq. (3) denotes that all advertisers are willing to purchase a guaranteed contract when its price is zero. It is worth mentioning that the above purchase behaviour assumptions are mainly for the general rational advertisers without budget constraints. The guaranteed contract price should not exceed the advertiser's value on an impression, which will be discussed in Section 3.3.

that they would pay for guaranteed contracts and RTB, and which cannot be observed in our dataset. We thus discuss the advertisers' strategic waiting behaviour as a future research topic in Section 6.

Consistent with the existing literature, we formulate $\theta(t_n, p(t_n))$ in the following form:

$$\theta(t_n, p(t_n)) = \exp \left\{ -\alpha p(t_n) \left(1 + \beta(t_N - t_n) \right) \right\}, \tag{4}$$

where α represents the price effect and β represents the time effect. This functional form is widely used in dynamic programming with various applications, such as selling flight tickets (Anjos, Cheng, & Currie, 2005), display advertising (Chen et al., 2014), and so on.

Based on Eq. (4), the demand for buying a guaranteed contract at time t_n can be computed as follows:

$$\eta(t_n) = \mathbb{I}_{\{t_n > 0\}} \sum_{i=0}^{n-1} f(t_i) \prod_{j=i}^{n-1} \left[1 - \theta(t_j, p(t_j)) \right] + f(t_n), \tag{5}$$

where $\mathbb{I}_{\{t\}}$ is an indicator function, $\sum_{i=0}^{n-1} f(t_i) \prod_{j=i}^{n-1} (1 - \theta(t_j, p(t_j)))$ computes the unfulfilled demand backlogged from the previous time periods and $f(t_n)$ is the expected number of advertisers arriving in the current time period.

3.2. RTB-based terminal value

Another distinguished feature of our model is that, the terminal value of this dynamic programming problem depends on the outcome of RTB. In RTB, impressions are usually sold separately through the seal-bid second-price auction (Ben-Zwi et al., 2015), and the existing literature has shown that such a mechanism enables truth-telling, namely, it is a weakly dominant strategy to bid at one's valuation (Narahari, 2014).

Let ξ be the number of advertisers who enter an RTB campaign, which can be interpreted as the competition level in RTB. By following the auction literature (Narahari, 2014), the expected

revenue from RTB can be obtained as follows:

$$\phi(\xi) = \int_{\Omega} x\xi(\xi - 1)g(x)[1 - F(x)][F(x)]^{\xi-2} dx, \tag{6}$$

where x is an advertiser’s bid, Ω is the range of bid, $g(\cdot)$ and $F(\cdot)$ are the density and cumulative distribution functions, respectively. Therefore, $\xi(\xi - 1)g(x)[1 - F(x)][F(x)]^{\xi-2}$ represents the probability that if an advertiser who bids at x is the second highest bidder, then one of $\xi - 1$ other advertisers must bid at least as much as he does and all of $\xi - 2$ other advertisers have to bid no more than he does. Usually, uniform or log-normal distributions are used for $g(\cdot)$ and $F(\cdot)$ to model the bid distribution (Narahari, 2014; Ostrovsky & Schwarz, 2011) and $\phi(\cdot)$ can be solved in closed-form if bids are uniformly distributed. However, neither distribution coincides with empirical data on many instances (Chen, 2016; Chen et al., 2014; Yuan, Wang, Chen et al., 2014). Thus, in this paper, $\phi(\cdot)$ will be learned from data and we will discuss this process in Section 4.

3.3. Censored upper bound for pricing

When making purchase decisions, an advertiser maximises his utility by comparing the expected costs from guaranteed contract and RTB. Due to the higher risk of RTB compared to PG, the guaranteed contract price should include a risk premium which measures the uncertainty or risk that the advertiser fails to win the RTB campaign. At time t_n , the censored upper bound of the guaranteed contract price can be characterised as follows:

$$\Phi(t_n) = \min \left\{ \underbrace{\phi(\xi(t_n)) + \delta(t_n)\psi(\xi(t_n))}_{:=\chi(t_n, \xi(t_n))}, \pi \right\}, \tag{7}$$

where π is the expected maximum value of an impression and $\chi(t_n, \xi(t_n))$ is the risk-aware upper bound – it is the sum of the expected payment in RTB and the risk premium. The risk premium is operationalised as the multiplication of the standard deviation $\psi(\xi(t_n))$ of payment prices in RTB and the advertiser’s risk preference $\delta(t_n)$. Similar to $\phi(\cdot)$, $\psi(\cdot)$ and π can be learned from data. And for $\delta(t)$, we model it with an exponential decay function $\delta(t_n) = \zeta e^{-\nu t_n}$ so its derivative $\delta(t)' \leq 0$. Here, ζ represents the degree of risk aversion and ν represents the time effect. Similar functional forms have been widely used in asset pricing and risk analysis literature (Wilmott, 2006). Note that, if N is large, $\delta(\cdot)$ ensures that $\chi(t_n, \xi(t_n))$ approaches $\phi(\xi(t_n))$ when the time is closer to the delivery day.

3.4. Revenue maximisation

We next formulate the publisher’s revenue maximisation problem. Let R be the publisher’s expected total revenue. It consists of the expected revenue from selling impressions through guaranteed contracts, denoted by R^{PG} , and the expected revenue from auctioning the remaining impressions in RTB, denoted by R^{RTB} . Thus, the revenue maximisation problem can be written as follows:

$$\max R = \left\{ \underbrace{\sum_{n=0}^N (1 - \omega\varpi) p(t_n)\theta(t_n, p(t_n))\eta(t_n)}_{:=R^{PG}} + \underbrace{\left[S - \sum_{n=0}^N \theta(t_n, p(t_n))\eta(t_n) \right] \phi(\xi(t_N))}_{:=R^{RTB}} \right\}, \tag{8}$$

$$\text{s.t. } 0 \leq p(t_n) \leq \Phi(t_n), \text{ for } n = 0, \dots, N, \tag{9}$$

$$0 \leq \sum_{n=0}^N \theta(t_n, p(t_n))\eta(t_n) \leq S, \tag{10}$$

where ω is the probability that the publisher fails to deliver a guaranteed impression, ϖ is the size of penalty proportional to the price so that the publisher needs to pay $\omega p(t_n)$ penalty if he fails to deliver a guaranteed impression which is sold at price $p(t_n)$, and the level of competition in RTB, $\xi(t_n)$, $n = 1, \dots, N$, can be measured by the average number of advertisers per impression as follows

$$\xi(t_n) = \frac{Q - \sum_{i=0}^n \theta(t_i, p(t_i))\eta(t_i)}{S - \sum_{i=0}^n \theta(t_i, p(t_i))\eta(t_i)}. \tag{11}$$

In the above revenue maximisation problem, ω , ϖ , Q , S , N are treated as model parameters so their values are assumed to be given when we use dynamic programming to solve the revenue maximisation problem.⁴ Eq. (9) specifies the boundaries of the guaranteed contract price at each time point, and Eq. (10) ensures the total amount of impressions sold via guaranteed contracts does not exceed the estimated total supply S in order to prevent over-selling. The decision variable $\mathbf{p} = [p(t_0), \dots, p(t_N)]$ is a vector of guaranteed contract prices which gives the pricing strategy and each \mathbf{p} has a corresponding allocation ratio γ representing the percentage of impressions that should be sold via guaranteed contracts. Therefore, the optimal pricing and allocation strategy can be denoted by (\mathbf{p}^*, γ^*) .

3.5. Optimal solution

The main challenge of solving the publisher’s revenue maximisation problem is that R^{RTB} is uncertain because the action of selling guaranteed contracts will affect both the supply and demand of impressions in RTB, which further impacts its expected revenue. Our problem is similar to the Knapsack problem (Kleinberg & Tardos, 2005), where the optimisation deals with a sequence of items with values and non-negative weights. In our problem, prices are affected by the allocation and they are further censored based on the buying behaviour of advertisers. As the expected total revenue can be divided into R^{PG} and R^{RTB} , when the remaining total supply and demand are given, R^{RTB} can be estimated consequently. So the optimal R^{PG} can be obtained using the recurrent structure of dynamic programming. Our solution is presented in Algorithms 1 and 2. Algorithm 1 initiates advertisers’ arrivals and computes the optimal total expected revenue based on different allocation schemes. For each allocation scheme, Algorithm 2 computes the optimal expected revenue of selling the guaranteed contracts. Fig. 2 presents a schematic view of our solution. In essence, we create a decision tree over time, in which each node represents a possible scheme of allocation and the corresponding price at a specific time point. The optimal allocation and pricing strategy is the trial over the tree which maximises the expected total revenue. Next, we explain a few key steps of both algorithms in detail.

As the demand for buying guaranteed contracts arrives following a Poisson process, together with Eq. (10), the upper bound u_n of total amount of sold impressions up to time t_n can be defined by $\min\{S, \sum_{i=0}^n f(t_i)\}$. Its lower bound l_n is given from time t_{n-1} and l_0 is zero. The possible sold impressions up to time t_n can

⁴ It should be noted that the advertiser pays when he buys a guaranteed contract and he can receive the penalty payment from the publisher if guaranteed impression is not delivered on the delivery day.

Algorithm 1 OPT-R.

```

1: Input:  $\alpha, \beta, \zeta, \eta, \omega, \kappa, \lambda, S, Q, T, N$ 
2:  $\mathbf{t} = [t_0, \dots, t_N]$ ; ▷ Initialisation
3: for  $n \leftarrow 0, \dots, N$  do
4:    $u_n \leftarrow \min\{S, \sum_{i=0}^n f(t_i)\}$ ; ▷ Sales upper bound
5:    $l_n \leftarrow l_{n-1} \mathbb{I}_{\{n>0\}}$ ; ▷ Sales lower bound
6:    $\mathcal{Y}_n \leftarrow \{l_n, l_n + 1, \dots, u_n\}$ ; ▷ Set of sales at time  $t_n$ 
7:   for  $y \in \mathcal{Y}_n$  do
8:      $\mathbf{H}_n(y) \leftarrow \text{OPT-H}(\alpha, \beta, \zeta, \eta, \omega, \kappa, \lambda, S, Q, T, n, y)$ ; ▷
9:   end for
10:  if  $n = N$  then
11:     $R(y) \leftarrow \mathbf{H}_n(y) + (S - y)\phi\left(\frac{Q-y}{S-y}\right)$ ;
12:  end if
13: end for
14: Output:  $R^* \leftarrow \max_{y \in \mathcal{Y}_N} \{R(y)\}$ ;  $\{\mathbf{H}_n(y) : n = 0, \dots, N; y \in \mathcal{Y}_n\}$ 

```

Algorithm 2 OPT-H.

```

1: Input:  $\alpha, \beta, \zeta, \eta, \omega, \kappa, \lambda, S, Q, T, n, y, \mathbf{H}_{n-1}$ 
2:  $\mathbf{z} \leftarrow$  an allocation matrix with size  $(y + 1) \times 2$ ; ▷ Initialisation
3:  $j \rightarrow 1$  ▷ Index initialisation of  $\tilde{\mathbf{H}}$ 
4: for  $k \leftarrow 1, \dots, (y + 1)$  do
5:    $p_{n,j,k} \leftarrow$  Eq. (12);
6:    $\Phi_{n,j,k} \leftarrow$  Eq. (7);
7:   if  $p_{n,j,k} > \Phi_{n,j,k}$  then ▷ Check price bound
8:     Continue;
9:   end if
10:   $\tilde{\mathbf{H}}_{n,j}(y) \leftarrow \mathbb{I}_{\{n>0\}} \mathbf{H}_{n-1}(\mathbf{z}(k, 1)) + p_{n,j,k} \mathbf{z}(k, 2)$ ;
11:   $j \rightarrow j + 1$ 
12: end for
13: Output:  $\mathbf{H}_n(y) \leftarrow \max_j \{\tilde{\mathbf{H}}_{n,j}(y)\}$ ;

```

then be generated, denoted by set \mathcal{Y}_n . For any $y \in \mathcal{Y}_n$, its optimal subset-sum can be computed by creating a $(y + 1) \times 2$ matrix \mathbf{z} , which contains all possible combinations of impressions sold up to time t_{n-1} and at time t_n , such that $\mathbf{z}(k, 1) + \mathbf{z}(k, 2) = y$, for $k = 1, \dots, y + 1$, and $\mathbf{z}(k, 1) = 0$ if $n = 0$. Since

$$\frac{\mathbf{z}(k, 2)}{\sum_{i=0}^n f(t_i) - \mathbf{z}(k, 1)} \propto \theta(t_n, p(t_n)),$$

the price of a guaranteed contract at time t_n can be obtained as

$$p_{n,j,k} = - \frac{\ln\{\mathbf{z}(k, 2)\} - \ln\left\{\sum_{i=0}^n f(t_i) - \mathbf{z}(k, 1)\right\}}{\alpha(1 + \beta(t_N - t_n))}. \tag{12}$$

We then obtain $\Phi_{n,j,k}$ by Eq. (7). If $p_{n,j,k} > \Phi_{n,j,k}$, the price $p_{n,j,k}$ will be removed from the solution space and this price will be replaced with the value calculated by the next index of k . If $p_{n,j,k} \leq \Phi_{n,j,k}$, we can calculate the corresponding revenue of selling guaranteed contracts up to time t_n , denoted by $\tilde{\mathbf{H}}_{n,j}(y)$. Then, the maximised revenue of selling guaranteed contracts up to time t_n is obtained, denoted by $\mathbf{H}_n(y)$. The iterations carry on until time t_N . For each $y \in \mathcal{Y}_N$, the expected total revenue can be obtained by adding $\mathbf{H}_n(y)$ with the corresponding expected revenue from RTB. Finally, the optimal solution can be obtained by comparing the expected total revenues from all candidates $\{R(y) : y \in \mathcal{Y}_N\}$ in the solution space. As the optimal selling amount and the corresponding price of each time step have been stored, we can use the index of the maximised revenue to obtain the optimal price vector \mathbf{p}^* and the corresponding allocation γ^* . Our solution is a greedy algorithm and the time complexity is $\mathcal{O}(NS^2)$, where \mathcal{O} is a notation used to classify algorithms according to how the running time grows as the input size

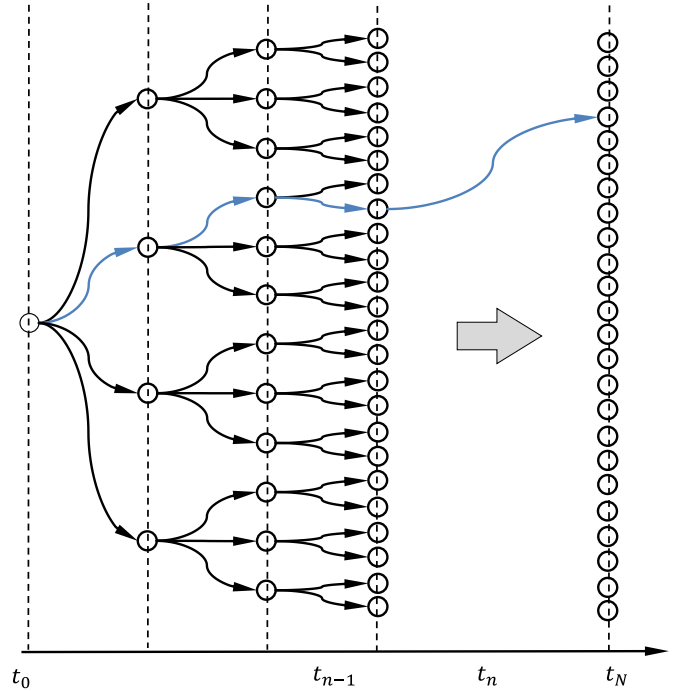


Fig. 2. Schematic view of the proposed optimal solution. Each node represents a possible scheme of allocation and pricing of impressions at the corresponding time steps.

grows. As N is much less than S , our solution is relatively scalable and efficient.

4. Data and experimental settings

We use an RTB dataset from a UK SSP to validate the proposed model of selling impressions via both guaranteed contracts and RTB.⁵ This dataset contains 1,378,971 RTB campaigns for 31 different ad slots over the period from 08 January 2013 to 14 February 2013. For each ad slot, RTB campaigns range from 7 to 20 continuous days. Therefore, we select the campaign records of 7 continuous days from all ad slots for experiments. Given an ad slot, the delivery day $[t_N, t_N]$ is randomly selected and the campaigns reported from this day is used for validation and testing (called the *test set*). The records of continuous 6 days prior to the delivery day are used to estimate the model parameters (called the *training set*). For those ad slots with only 7 days data, the 7th day is set to be the delivery day. In our dataset, all bids in RTB campaigns are quoted in terms of cost-per-mille (CPM), which is a measurement corresponds to the value of 1000 impressions (Yuan, Wang, Li et al., 2014).

The total supply S and demand Q of impressions for a specific ad slot in the delivery day can be predicted by time series or regression models. As our primary intention here is not to discuss a prediction model, Q and S are simply given by the test set. Fig. 3 summarises the total supply and demand of impressions for ad slots in both training and test sets. There are 5 slots excluded from the original dataset because the competition level of the slot in RTB is less than 2 in the training set (i.e., $\xi < 2$). In such cases, advertisers are able to obtain the needed impressions in RTB at a very low price or even the reserve price. Therefore, it is very unlikely for these advertisers to buy any guaranteed contracts in advance. Overall, it appears that the total supply levels are similar

⁵ The dataset has also been used in several recent display advertising studies (Chen, 2016; Chen et al., 2014; Yuan, Wang, Chen et al., 2014).




Fig. 3. Summary of the RTB dataset: (a) the total demand; (b) the total supply; (c) the average per impression competition level.

Table 1

Summary of clustered ad slots, where numbers in round brackets are standard deviations.

Group	1		2	
	Number of ad slots		20	
Set	Training	Test	Training	Test
Payment price	0.98 (0.09)	0.99 (0.08)	0.73 (0.46)	0.56 (0.36)
Winning bid	1.13 (0.17)	1.1 (0.1)	2.32 (1.17)	1.84 (1.04)
ξ	8.92 (3.24)	8.15 (1.18)	3.39 (0.59)	3.51 (0.81)
Ratio of payment price to winning bid	88.95% (4.54%)	92.88% (2.15%)	32.2% (9.9%)	37.18% (10.58%)

across ad slots while the total demand levels are significantly different. It is worth noting that, within a day, the competition level ξ varies significantly from its mean as many advertisers join RTB at peak hours between 6 a.m. and 10 a.m., making ξ on average 118.96% higher than other hours. In the experiments, we assume guaranteed contracts can be sold 31 days prior to the delivery day, and as discussed in Section 3, the advertiser arrivals follow a Poisson process during the 31 days. To simplify the discussion and without loss of generality, intensity λ is set as a constant. Specifically, 20% of Q is assumed to arrive at time 0 and another 20% of Q is considered to arrive prior to the delivery day. Therefore, $\lambda = 0.2Q/30$.⁶

We divide 26 ad slots into two groups using the K-Means clustering (Bishop, 2006) based on the competition level ξ in the training set. Table 1 presents a brief summary of the major statistics of both groups. The table shows that, group 1 with a high competition level has the average 8.92 (in the training set) and 8.15 (in test set) advertisers bidding per RTB auction, much higher than those of group 2 with a low competition level. Consistent with Eq. (6), the average of per-auction payments of group 1 are also higher than in group 2. Interestingly, group 2 has a higher average winning bid. By further investigating the distributions of winning bids and payment prices, we find that the counter-intuitive result is due to the different price patterns over two groups, which is demonstrated in Fig. 4. Fig. 4 uses two examples, slot 24 in group 1 and slot 15 in group 2. For slot 24, the distributions of winning bids and payment prices are bell shaped, and payment price has a higher peak and a slightly smaller mean. Differently, for ad slot 15, the distributions of winning bids and payment prices are similar to Gaussian mixture distributions (Bishop, 2006), and payment prices are much lower than winning bids. The ratio of payment price to winning bid for group 1 is 88.95% while it is 32.2% for group 2. This suggests that, for ad slots in the group with a low competition

level, there is greater potential to optimise the selling mechanism to increase the revenue. Further comparing the training set and the test set in both group 1 and 2, we observe from Table 1 that most statistics are close and consistent. Thus, it makes sense that we use the training set to estimate the model parameters, and validate the model in the test set.

We further use the dataset to estimate parameters $\phi(\xi)$, $\varphi(\xi)$ and $\chi(t, \xi(t))$ when ξ is given, which are visualised in Fig. 5. Following existing literature (Chen, 2016; Chen et al., 2014), we implement the locally weighted regression scatterplot smoothing (LOWESS), polynomial regression, and sigmod methods.⁷ The LOWESS method combines multiple weighted polynomial regression models (Cleveland, 1979) and we follow the implementation settings of Algorithm 1 given in Chen et al. (2014). The sigmod method uses a scaled sigmod function to approximate the training data, where the scaling parameters are calibrated based on the root mean squared errors. Finally, the LOWESS method is selected for parameters estimation as it fits the training data best. Given time t , $\delta(t)$ is calculated and then $\chi(t, \xi(t))$ is obtained. We set π as the maximum value of average bids (hourly basis) in the training set, then a surface of the price upper bound $\Phi(t)$ is obtained.

5. Results

In this section, we first present our experimental results on the model's performance, and then investigate the model's robustness by considering: (i) the effect of uncertainty in supply and demand of impressions; and (ii) advertisers with different valuations. The robustness analysis can provide insights for daily operations as the two situations may occur in practice.⁸

5.1. Optimal pricing and allocation

To illustrate how the proposed model can optimally price and allocate impressions between guaranteed and spot markets, Figs. 6 and 7 present examples of ad slot 24 from group 1 and ad slot 15 from group 2. We deliberately show the results of these two slots because they are typical instances in the corresponding group. Also, as the distributions of winning bids and payment prices of both slots in RTB have been shown in Fig. 4, we can clearly see if the optimal prices of guaranteed contracts are feasible.

Fig. 6 illustrates the optimal pricing strategies suggested by the model as well as the effects of the advertisers' risk preference (i.e., ζ and ν) on the optimal prices of guaranteed contract. Comparing the sub-figures in the same column, we can see that there is a general trend of price increase over time for both slots. The price on

⁷ Other statistical or machine learning methods may also be used here, however, discussing the best prediction method is out of the scope of this paper.

⁸ We thank one of the reviewers for suggesting the extensions in Sections 5.2 and 5.3.

⁶ Our results with different values for λ show that, consistent with intuition, smaller (larger) λ leads to less (more) guaranteed contracts being sold out. However, the main insights from the analysis remain the same.

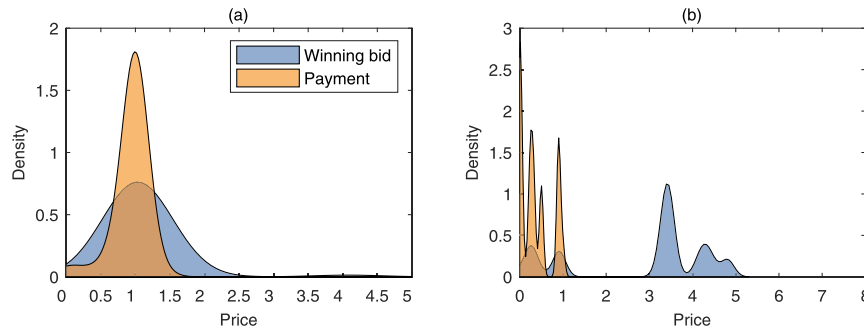


Fig. 4. Examples of distributions of winning bids and payment prices in the training set: (a) ad slot 24 in group 1; (b) ad slot 15 in group 2.

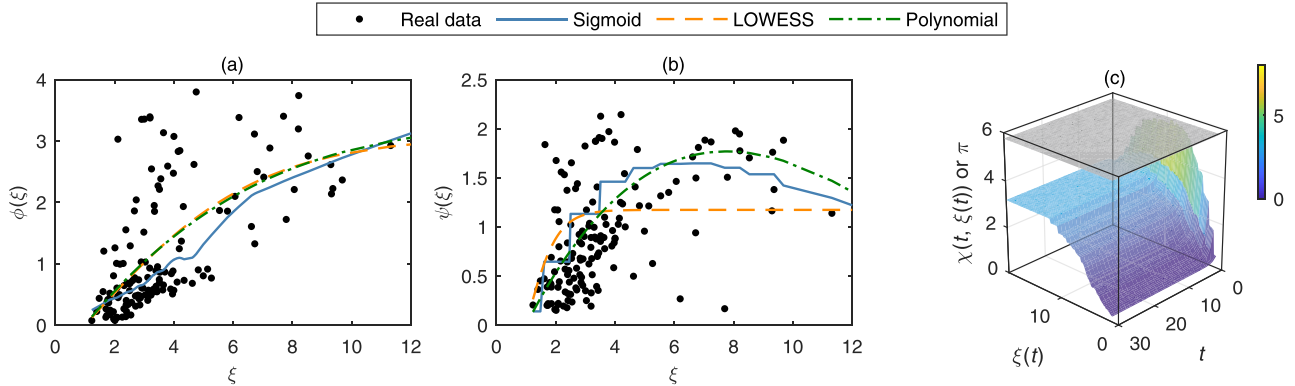


Fig. 5. Examples of estimating the censored price upper bound of guaranteed contracts for ad slot 6 in group 2 based the training data (hourly basis): (a) $\phi(\xi)$; (b) $\varphi(\xi)$; and (c) surfaces of $\chi(t, \xi(t))$ and π (in grey colour).

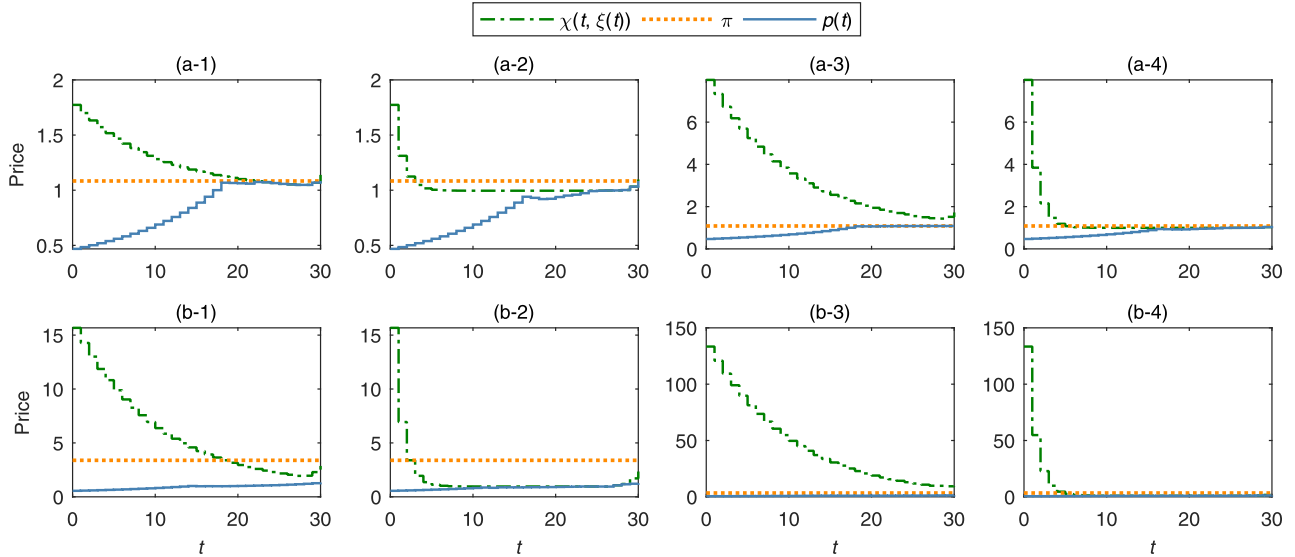


Fig. 6. Examples of optimal pricing of guaranteed contracts suggested by the model for: (a) ad slot 24 in group 1 and (b) ad slot 15 in group 2. Parameters are set differently in the subplots: (1) $\zeta = 10, \nu = 0.1$; (2) $\zeta = 10, \nu = 0.9$; (3) $\zeta = 90, \nu = 0.1$; (4) $\zeta = 90, \nu = 0.9$.

slot 15 increases more steadily. This is because there is less demand for buying guaranteed contracts prior to the delivery day due to the low competition level. Also, as less advertisers arrive, the price growth is not significant over time. On the contrary, as it is difficult to obtain impressions in RTB for slot 24 with high competition level, more advertisers would be willing to secure the impressions in advance through guaranteed contracts. Consequently, the guaranteed contracts can be sold at higher prices. This is con-

firmed by the observation in Fig. 6 that the price increases steeply and then be censored by the upper bound $\Phi(t)$, which is the minimum value between π and $\chi(t, \xi(t))$, as defined in Eq. (7). Furthermore, by comparing the sub-figures in the same row, Fig. 6 shows how the advertisers' risk preference parameters ζ and ν affect $\chi(t, \xi(t))$ and the optimal price. The value of $\chi(t, \xi(t))$ increases significantly with the increase of ζ . Since $\chi(t, \xi(t))$ is an exponential decay function of ν , it converges quickly to $\phi(\xi(t))$ if

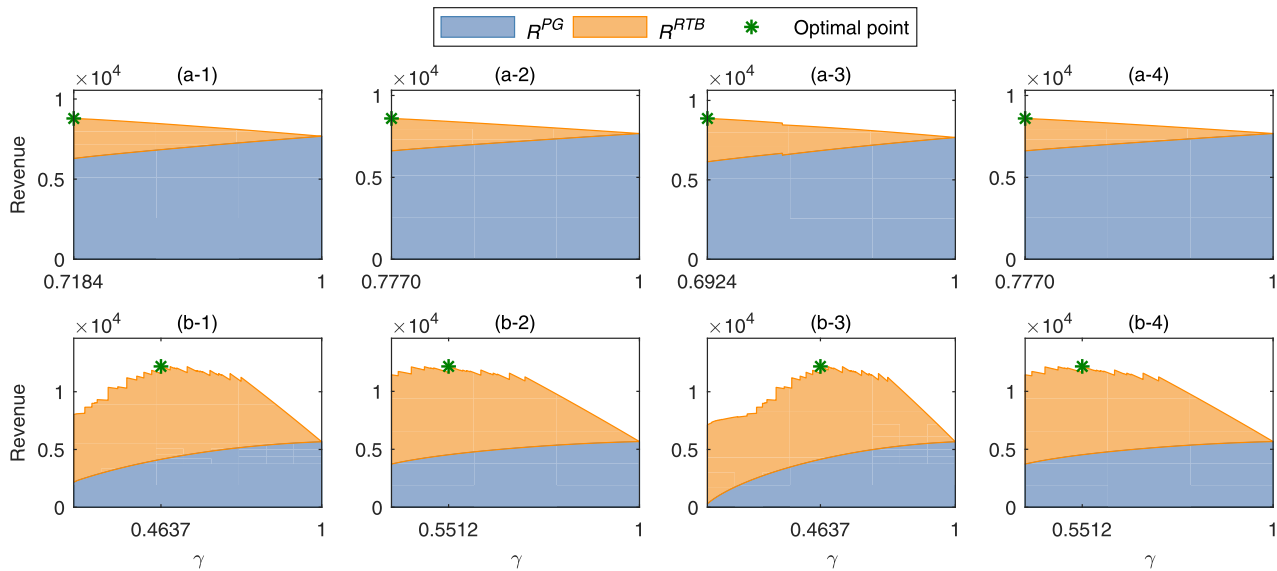


Fig. 7. Examples of optimal allocation of guaranteed contracts suggested by the model for: (a) ad slot 24 in group 1 and (b) ad slot 15 in group 2. Parameters are set differently in the subplots: (1) $\zeta = 10, \nu = 0.1$; (2) $\zeta = 10, \nu = 0.9$; (3) $\zeta = 90, \nu = 0.1$; (4) $\zeta = 90, \nu = 0.9$.

ν increases. The optimal price is more sensitive to ν . However, the impacts of ν and ζ on the optimal price is mild.

We next investigate the optimal allocation in Fig. 7. This figure shows the details of the corresponding optimal allocations made by the model for the same ad slot and under the same experimental settings of Fig. 6. Obviously, the model suggests different allocation strategies for different slots. For slot 24 with high competition level, the model allocates more future impressions into guaranteed contracts, and they constitute a large percentage of the expected total revenue. For slot 15 with low competition level, the model allocates only a small amount of impressions in advance. As a result, the average payment price of impressions auctioned in RTB increases significantly, and the expected total revenue is largely contributed by RTB rather than PG. Furthermore, Fig. 7 also indicates the effects of advertisers' risk preference, ζ and ν , on optimal allocation. First, the proportion of impressions allocated to PG, γ , is decreasing in ζ . This is because guaranteed contract prices are slightly higher if ζ is large so less advertisers would be willing to buy guaranteed contracts. Second, γ is increasing in ν . This is because if ν increases: (i) advertisers are more risk averse and more sensitive to time; (ii) $\chi(t, \xi(t))$ decreases quickly and the guaranteed prices will be slightly lower so more advertisers would be willing to buy in advance.⁹

The overall results of the model performance on ad slots of two clustered groups are summarised in Fig. 8. We compare the expected total revenue given by the model with both the expected RTB revenue and the actual RTB revenue. The expected RTB revenue is calculated by multiplying the average per-auction payment with the total number of impressions, while the actual RTB revenue is the total revenue reported by actual RTB campaigns in the test set.¹⁰ Our results show that selling impressions via both PG and RTB indeed increase the total revenue for two groups with different competition levels. With the increase of ζ , the increase in revenue converges to its maximum amount. Furthermore, by comparing the allocation for group 1 and 2, more impressions are allo-

cated into guaranteed contracts in group 1 when the competition level is high, consequently, the expected total revenue is largely contributed by guaranteed contracts. However, as mentioned previously in Table 1, there is greater potential to improve revenue in ad slots with low competition level, as its average ratio of payment price to winning bid is only around 30%, compared to 90% for the ad slots with high competition level. Therefore, if the ratio of payment to valuation¹¹ increases for the ad slots with low competition level, the revenue increase will be more significant, and our numerical analysis shows that in some cases the expected total revenue can even be doubled.

5.2. Incorporating uncertainty into model updating

We next further investigate the model by considering the effect of uncertainty in supply and demand. In the model setting, given a specific ad slot, S and Q are assumed to be well estimated by the forecasting models at the current time t_0 . However, they may change over time due to uncertainty.¹² As is shown in Eq. (11), the uncertainty from both the supply or demand will finally impact the model performance through ξ the per-impression competition level in RTB. For model tractability and ease of exposition, we can fix S and just consider the random behaviour of Q over time.¹³ Let Q_{n+1} be the rest of total demand in $[t_n, t_{\bar{N}}]$ estimated at time t_{n+1} and let Q_n be the rest of total demand estimated at

¹¹ The ratio of payment to valuation is equivalent to the ratio of payment to winning bid in Table 1.

¹² If we assume they move either up and down separately, their behaviour over time can be captured by two multi-step binomial tree or lattice models (Chen et al., 2014; Cox, Ross, & Rubinstein, 1979) and there will be four different combinations of their estimations at each time step. If they can either move upwards, downwards, or stay unchanged in a given time period, their behaviour can be modelled by two multi-step trinomial tree or lattice models (Boyle, 1986) and there will be eight different combinations of their estimations at each time step.

¹³ The supply of impressions are usually stable with relatively low uncertainty, as the traffic or the number of visitors to the web pages are usually predictable. For example, tools like Google Analytics (<https://analytics.google.com>) and SimilarWeb (<https://www.similarweb.com>) offer detailed website traffic estimation service. Lee and Leckenby (1999) and Ilfeld and Winer (2002) have also shown that many popular websites have rather stable traffic over a short-term period such as week or month.

⁹ The optimal allocations are same in Fig. 7 (a-2) and (a-4) as well as in Fig. 7 (b-2) and (b-4). This because the effects of ζ and ν are not significant in the chosen ad slots under our experimental settings. In Appendix B, we provide two additional examples to further justify our analysis.

¹⁰ This is also called *ground truth* in machine learning literature.

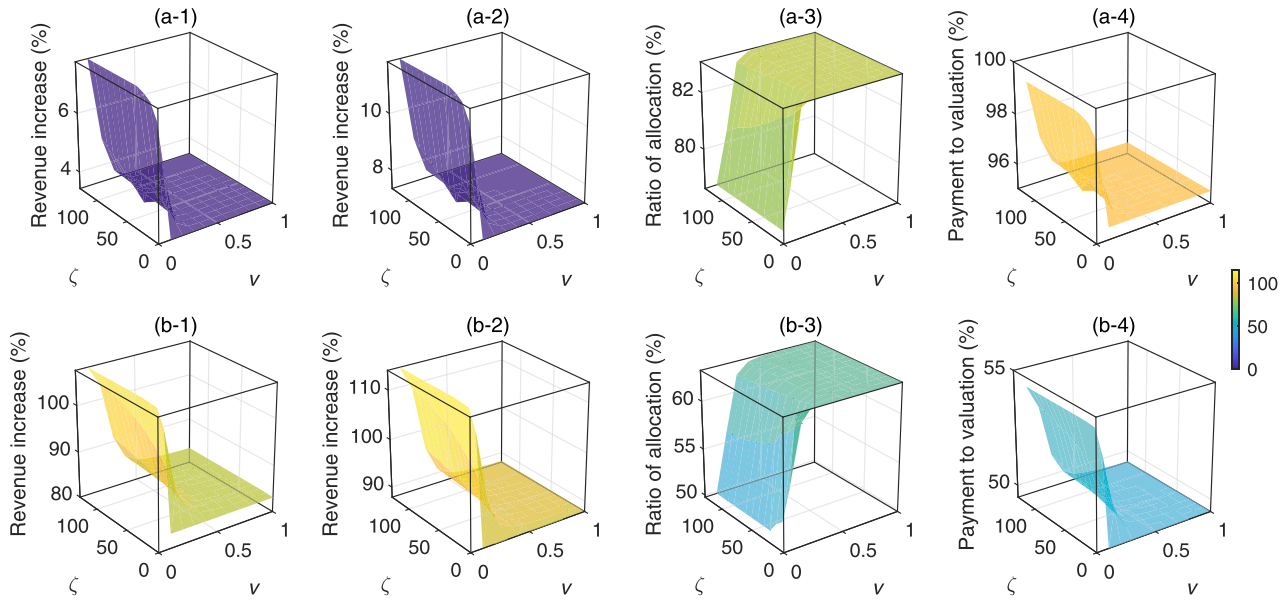


Fig. 8. Overall results of the model performance on ad slots in: (a) group 1; (b) group 2. The sub-plots show: (1) the average revenue increase of the model to the expected RTB; (2) the average revenue increase of the model to the actual RTB; (3) the average ratio of selling impressions in advance made by the model; (4) the average ratio of payment to valuation made by the model.




Fig. 9. Examples of optimal pricing of guaranteed contracts suggested by the model by adding 10% uncertainty level ($\epsilon = 0.1$) for: (a) ad slot 24 in group 1 and (b) ad slot 15 in group 2. Parameters $\zeta = 10, \nu = 0.1$.

time t_n . We can simply use a random walk structure to update uncertainty in Q as follows: $Q_{n+1} \leftarrow Q_n(1 + \epsilon\varepsilon)$, where ϵ is the uncertainty level expressed as a percentage and ε is a white noise such that $\mathbb{E}[\varepsilon] = 0$ and $\text{var}[\varepsilon] = 1$. In experiments, we set $\epsilon = 10\%$. When the time step moves, the sold impressions are then removed from S and Q accordingly. We then generate a new estimate for the rest of demand, update the corresponding model parameters, and recompute the optimal prices and allocations of guaranteed contracts for the rest of the days in the selling period. The iterations will continue until time t_N . By updating the uncertainty, the time complexity becomes $\mathcal{O}(N^3S^2)$.

Fig. 9 compares the optimal prices with and without considering the uncertainty for ad slot 24 in group 1 and ad slot 15 in group 2. The results show that uncertainty affects the optimal prices marginally. The increasing trend of the price over time remains the same, although there are small fluctuations in price movement. Table 2 summarises how uncertainty impacts the optimal allocation and the expected revenue. The results show that that uncertainty indeed impacts the optimal impression allocation between PG and RTB and their respective expected revenues. Such impacts are marginal on the expected revenue, but can change the allocation significantly.

5.3. Advertiser segmentation

Advertisers who bid for the same ad slot can have different valuations on the impressions (Abraham, Athey, Babaioff, & Grubb, 2013; Sayedi, 2018). As discussed in Section 3.2, truth-telling is a weakly dominant strategy in RTB. Thus we can segment advertisers with different valuations based on their bids. For each ad slot, we use the K-Means clustering to divide the RTB campaigns into two subgroups: subgroup 1 represents impressions bid by advertisers who have high valuations while subgroup 2 represents impressions with low-value advertisers. Table 3 presents a brief summary of the major statistics of subgroups of all 26 ad slots. Similar to Table 1, data expresses similar patterns in both training and test sets so that we can use the training set to develop prediction and pricing models for the future impressions in the test set. In group 1 (where slots have with a high competition level), subgroup 1 has the average 14 advertisers bidding per RTB auction, much higher than that of subgroup 2. In group 2 (where slots have a low competition level), subgroup 1's competition level is 4 or 5, almost double subgroup 2. The winning bid and the payment price of subgroup 1 are all higher than those of subgroup 2 because they are positively correlated with the competition level. One interesting finding is that, in group 1, the ratio of payment price to win-

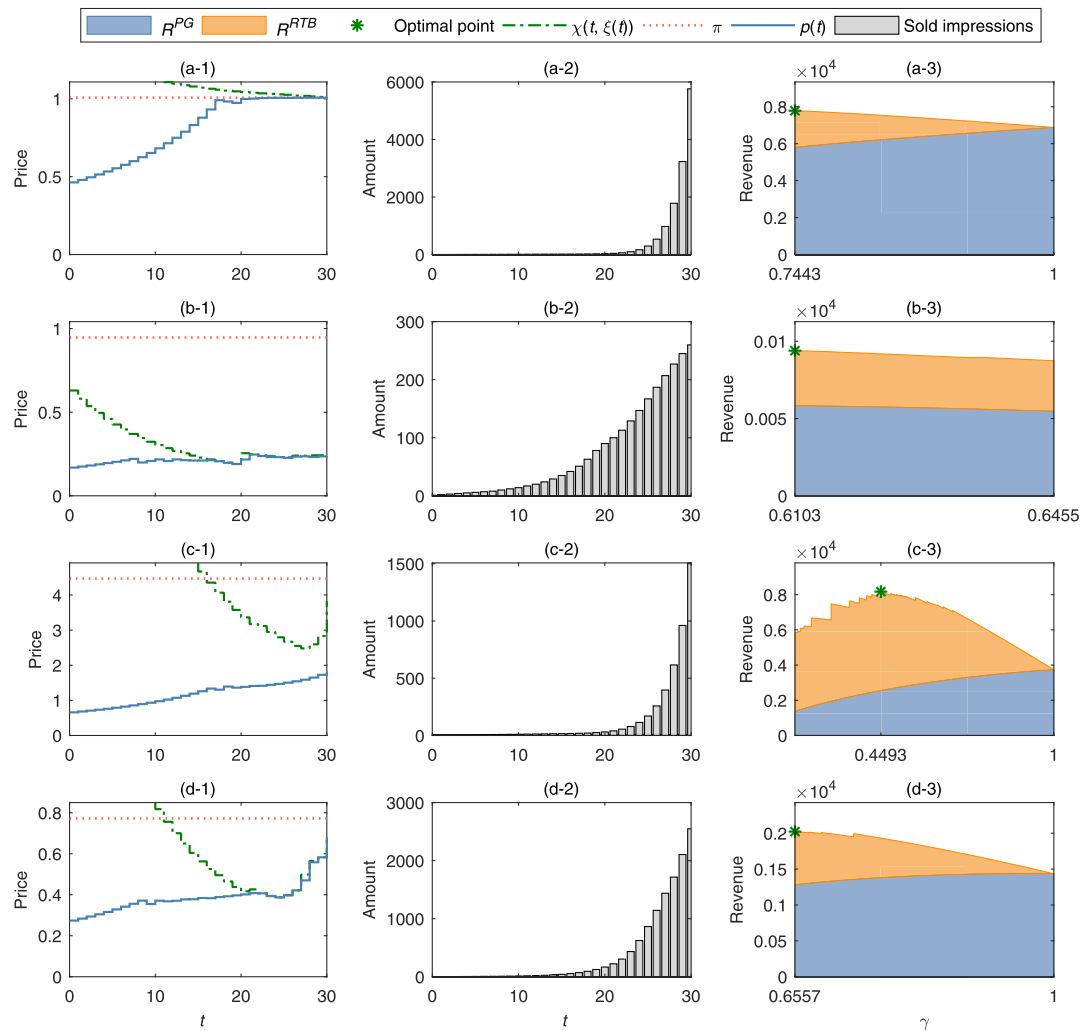


Fig. 10. Examples of optimal pricing, allocation and selling of guaranteed contracts suggested by the model. Sub-plots: (a) subgroup 1 of ad slot 3 in group 1; (b) subgroup 2 of ad slot 3 in group 1; (c) subgroup 1 of ad slot 15 in group 2; and (d) subgroup 2 of ad slot 15 in group 2. Sub-plots: (1) guaranteed prices and the boundaries over time; (2) sold guaranteed impressions over time; (3) expected total revenue and the optimal allocation. Parameters are set $\zeta = 10, \nu = 0.1$.

Table 2

Summary of changes of model allocation and revenue for clustered ad slots, where numbers in round brackets are standard deviations, where $\epsilon = 10\%$.

Group	Setting		Changes (%)			
	ζ	ν	γ	R^{RTB^a}	R^{PG}	R
1	10	0.1	8.29 (4.44)	-19.59 (6.16)	6.64 (4.28)	-1.15 (0.46)
	10	0.9	2.54 (2.97)	-6.84 (10.27)	3.16 (3.43)	0.38 (0.3)
	90	0.1	37.72 (19.21)	-41.77 (6.79)	23.88 (12.2)	-3.58 (0.55)
2	90	0.9	2.54 (2.97)	-6.84 (10.27)	3.16 (3.43)	0.38 (0.3)
	10	0.1	-8.74 (13.94)	-4.33 (23.76)	-4.3 (9.4)	-3.97 (14.44)
	10	0.9	-6.71 (15.18)	-1.75 (17.97)	-4.59 (9.38)	-1.81 (8.96)
	90	0.1	-12.72 (16.71)	-4.03 (24.15)	-6.05 (10.4)	-4.94 (13.95)
	90	0.9	-6.71 (15.19)	-1.74 (17.99)	-4.55 (9.39)	-1.79 (8.96)

^a Ad slot 26 in group 1 is excluded in the computation as the model suggests all impressions to be sold in advance via guaranteed contracts so its RTB revenue is 0.

ning bid in subgroup 1 is close to subgroup 2, all around 60%. This is because the winning advertisers offer high bids in most of RTB campaigns of ad slots 23–25 (whose ratios are around 30%) though the rest three slots in group 1 have the ratios around 95%.

Fig. 10 presents examples of optimal pricing, allocation and selling of guaranteed contracts for the subgroups of ad slots 3 and 15.

We use ad slot 3 to replace ad slot 24 to represent group 1 because impressions from the latter are all suggested to be sold via guaranteed contracts by our model so it is impossible for visualising the optimal allocation. Although small price fluctuations can be seen in some cases, there is a general trend of price increase over time in both subgroups in both markets. Impressions with different

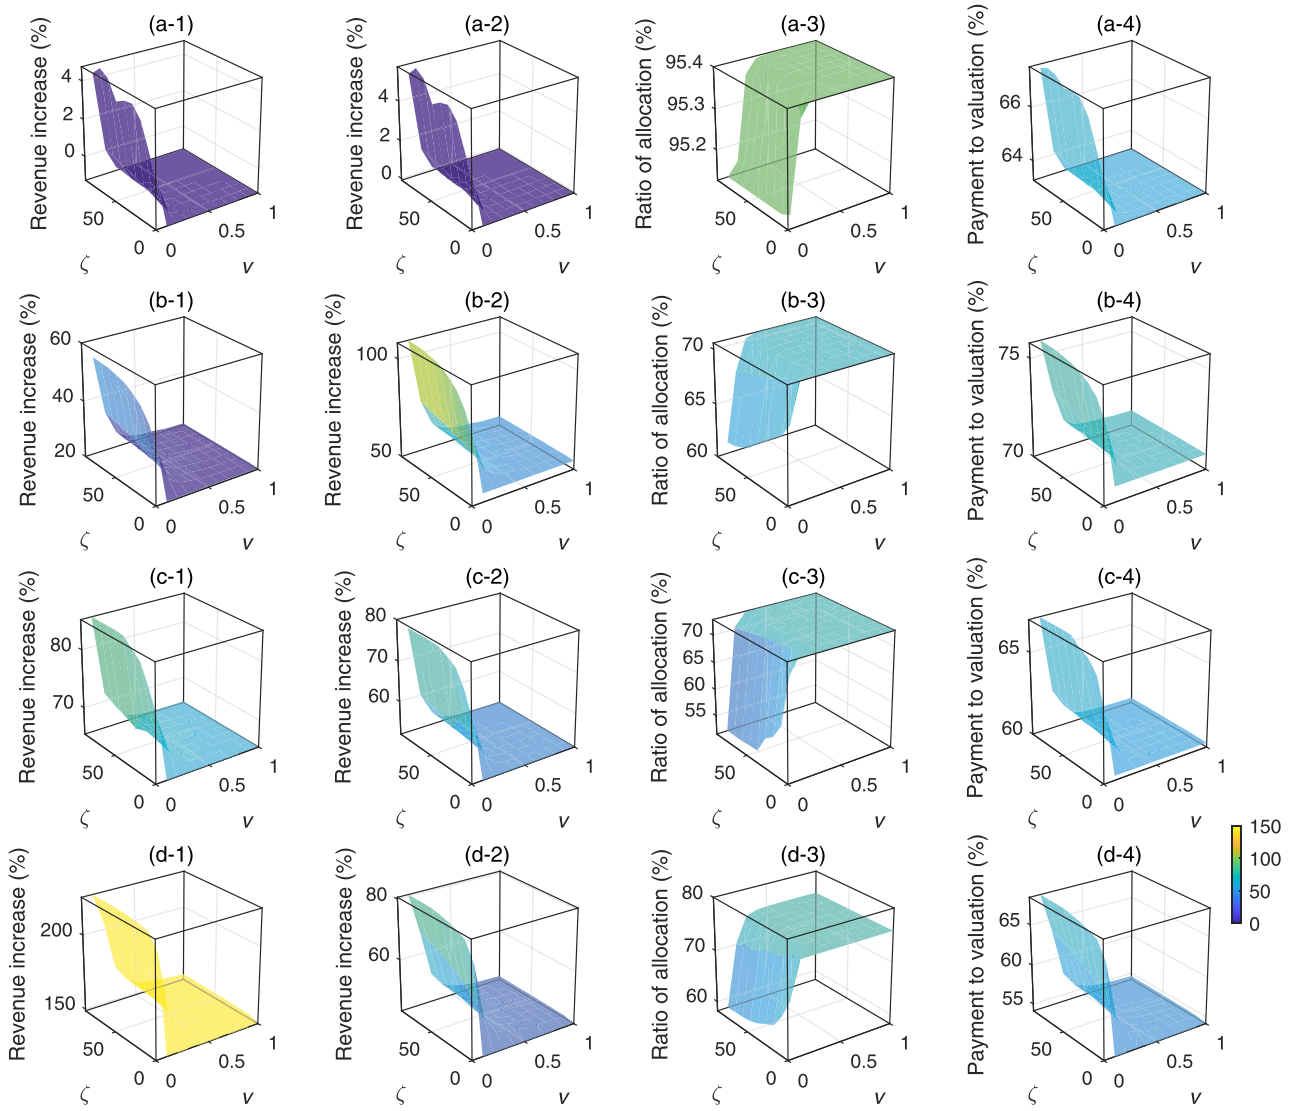


Fig. 11. Overall results of the model performance on ad slots in: (a) subgroup 1 of group 1; (b) subgroup 2 of group 1; (c) subgroup 1 of group 2; (d) subgroup 2 of group 2. The sub-plots show: (1) the average revenue increase of the model to the expected RTB; (2) the average revenue increase of the model to the actual RTB; (3) the average ratio of selling impressions in advance made by the model; (4) the average ratio of payment to valuation made by the model.

Table 3
Summary of clustered subgroups of ad slots, where numbers in round brackets are standard deviations.

		Set			
		Training		Test	
		1	2	1	2
Payment price	1	1.17 (0.12)	0.65 (0.34)	1.16 (0.12)	0.62 (0.4)
	2	1.09 (0.61)	0.16 (0.07)	1.03 (0.62)	0.2 (0.09)
Winning bid	1	2.63 (1.45)	1.0 (0.04)	2.62 (1.49)	0.99 (0.03)
	2	3.25 (1.48)	0.62 (0.22)	3.36 (2.03)	0.53 (0.2)
ξ	1	14.39 (5.26)	5.94 (2.7)	14.86 (6.39)	5.53 (2.91)
	2	4.1 (1.04)	2.61 (0.48)	5.17 (2.25)	2.71 (0.74)
Ratio of payment price to winning bid	1	61.98% (31.56%)	63.16% (32.63%)	63.69% (34.23%)	61.59% (39.75%)
	2	38.15% (19.19%)	28.4% (10.33%)	42.51% (25.96%)	35.07% (11.69%)

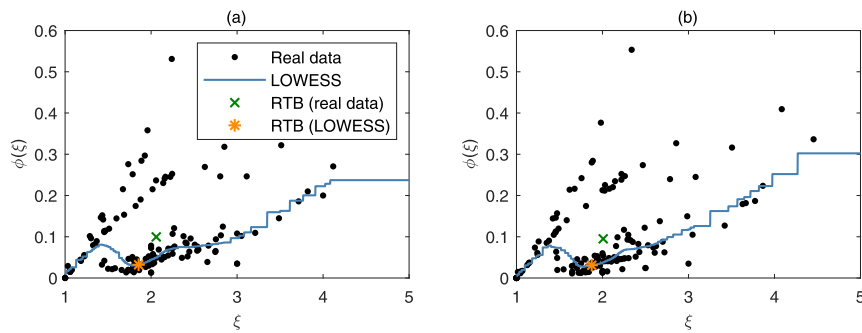


Fig. 12. Examples of $\phi(\xi)$ estimation for subgroup 2 of: (a) ad slot 14; (b) ad slot 19.

values are sold at totally different prices. It is worth emphasising that, in group 1, more high-valued impressions are encouraged to be sold in advance and the total revenue is mainly contributed by guaranteed contracts. This is not only because of the high competition level but also because advertisers with high valuations have great potential to accept a high payment price. Therefore, as shown in Fig. 10(a-1), price increases quickly and then be censored by the upper bound.

Fig. 11 summarises the overall results of the model performance on the subgroups of all 26 ad slots. Fig. 11 (a-3) shows the average γ of subgroup 1 in group 1 is about 95%, further confirming the replacement of ad slot 24 in Fig. 10. The performance metrics of the model are close for the rest of cases: (i) low-valued impressions in group 1; (ii) high-valued impressions in group 2; (iii) low-valued impressions in group 2. The extreme high revenue increase in Fig. 11 (d-1) is due to several outliers and we further explain the reason in Fig. 12. We use the LOWESS method to fit the $\xi - \phi$ relationship and use the fitted ϕ to compute the expected RTB revenue for benchmark as well as for optimisation. As shown in Fig. 12, the estimated ϕ is smaller than the average payment in the real data. Therefore, the expected RTB revenue is smaller than the actual RTB revenue, which then gives extremely high revenue increases on several ad slots.

6. Conclusion

In this paper, we study a novel approach of selling display ad impressions via both guaranteed contracts and RTB. We take into account the buying behaviour of advertisers, i.e., risk aversion and stochastic arrivals, and employ dynamic programming to find an optimal pricing and allocation strategy that maximises the publisher's expected total revenue. Our optimal solution is derived based on the solution frame to the Knapsack problem, and is relatively scalable and efficient. We further validate the results using a commercial RTB dataset. The experimental results show that, selling via both guaranteed contracts and RTB can significantly improve the publisher's total revenue. Furthermore, for impressions from ad slots with a high competition level, a large percentage of future impressions should be sold in advance via guaranteed contract, and for impressions from ad slots with a low competition level, the revenue is largely collected by RTB. However, their revenue increases are more significant. This is mainly due to the fact that advertisers pay much less than their valuations in RTB which gives more margins for PG to increase the revenue. In addition, the experimental results show that our model is robust under supply or demand uncertainty and when advertisers have different valuations on the impression.

This study has some limitations, which can be addressed as future research directions. First, due to the complex nature of combining both channels, we focus on a simplified model setup in the paper, while keeping the major features of guaranteed contracts and RTB. Our model setup can be relaxed to the case when there are multiple separate groups of impressions with different delivery time periods. For instance, the impressions from one ad slot on a specific day can be treated as a separate group, then for each group, our model can be applied. Therefore, an advertiser can buy impressions from the same ad slot with different delivery time separately. Second, we discuss a simple standardised guaranteed contract, in which the advertiser has no further control of advertising delivery in the delivery period. This setting is consistent with the exiting literature as reviewed in Section 2. It is possible that an advertiser would like to gain further control of advertising, e.g., deciding when to advertise in the future delivery period. Flexible guaranteed contracts like ad options can be the suitable mechanisms to use (Chen & Kankanhalli, 2019; Chen & Wang, 2015; Chen, Wang, Cox, & Kankanhalli, 2015). However, in these studies, the authors only investigate the contract pricing but not both allocation and pricing. Also, the ad option pricing models discussed in those studies are not optimal. Developing a new revenue maximisation model for optimal pricing and allocation of ad options can be an interesting future topic. Third, when guaranteed contracts are sold, how to prioritise the contracts in the delivery period is not discussed in our paper. This is a different problem that is out of the scope of our study because our model focuses on the sales framework (i.e., the optimal allocation of the future inventories into two channels and the corresponding optimal guaranteed contract prices). In the delivery period, the publisher can give equal priority to the guaranteed contracts or prioritise some contracts based on the pre-sold guaranteed contract price or other metrics (Feldman et al., 2009; Ghosh et al., 2009). Therefore, another interesting future direction is developing a unified model which integrates guaranteed contract selling and premium impressions delivery operation. Finally, our model simplifies the advertiser's behaviour and decision by using a ratio that represents the proportion of those who want to buy an impression, and we do not explicitly model the advertisers' strategic behaviour of deliberation over PG and RTB. Future research can model the strategic behaviour of advertisers using a utility function (Aviv & Pazgal, 2008; Su, 2007), and study how such a strategic behaviour impacts the publisher's optimal pricing and allocation decisions.

Appendix A. Notations

This appendix is not an essential part of the paper but may help readers reference the key notations used throughout the paper.

Table 4
Summary of the key notations.

Notation	Description
$t_0, \dots, t_N, t_{\bar{N}}$	Discrete time points: $[t_0, t_N]$ is the period to sell the guaranteed contracts; $[t_N, t_{\bar{N}}]$ is the period that the impressions are created, auctioned off (in RTB) and delivered.
Q	Estimated total demand for impressions in $[t_N, t_{\bar{N}}]$.
S	Estimated total supply of impressions in $[t_N, t_{\bar{N}}]$.
R^{PG}	Expected revenue from selling guaranteed contracts.
R^{RTB}	Expected revenue from selling the remaining impressions in RTB.
$p(t_n)$	Price of a guaranteed contract sold at time $t_n, n = 0, \dots, N$.
γ	Ratio of allocation of impressions into guaranteed contracts.
ω	Probability that the publisher fails to deliver a guaranteed impression in $[t_N, t_{\bar{N}}]$.
ϖ	Size of penalty so the publisher needs to pay $\varpi p(t_n)$ penalty if he fails to deliver a guaranteed impression which is sold at $p(t_n)$.
$\eta(t_n)$	Total (accumulative) arrived but unfulfilled demand at time t_n .
$\theta(t_n, p(t_n))$	Proportion of advertisers who are willing to buy an impression in advance at time t_n and at price $p(t_n)$.
α	Price effect in $\theta(t_n, p(t_n))$.
β	Time effect in $\theta(t_n, p(t_n))$.
$\delta(t_n)$	Risk preference for a buyer at time t_n .
ζ	Risk level in $\delta(t_n)$.
ν	Time effect in $\delta(t_n)$.
$\xi(t_N)$	Per-auction competition level of RTB in $[t_N, t_{\bar{N}}]$.
$\phi(\xi(t_n))$	Expected payment from an impression in RTB for the given $\xi(t_n)$.
$\psi(\xi(t_n))$	Expected risk of an impression in RTB for the given $\xi(t_n)$.
$\chi(t_n, \xi(t_n))$	Risk-aware upper bound of the guaranteed contract price at time t_n .
$\Phi(t_n)$	Censored upper bound of the guaranteed contract price at time t_n .
π	Expected maximum value on an impression.
λ	Intensity of the Poisson process describing the demand arrivals.
$f(t_n)$	Expected number of arrivals in $[t_{n-1}, t_n]$.
\mathcal{O}	Execution time required by an algorithm.
$\mathbb{E}[\cdot]$	Expectation.
$\text{var}[\cdot]$	Variance.
$g(\cdot)$	Density function of advertiser's bid.
$F(\cdot)$	Cumulative distribution function of advertiser's bid.
u_n	Sales upper bound at time t_n in Algorithm 1.
l_n	Sales lower bound at time t_n in Algorithm 1.
\mathcal{Y}_n	Set of sales at time t_n in Algorithm 1.
$\mathbf{H}_n(y)$	R^{PG} when selling y guaranteed contracts up to time t_n in Algorithm 2 $y \in \mathcal{Y}_n$.
$\mathbf{H}_n(y)$	Optimal R^{PG} in Algorithms 1 and 2, $y \in \mathcal{Y}_n$.
$R(y)$	Total revenue of selling y guaranteed contracts up to time t_n in Algorithm 1, $y \in \mathcal{Y}_n$.

Appendix B. Examples of optimal pricing and allocation

Here, we provide two additional examples to justify our analysis in Section 5.1. Fig. 13 shows examples from ad slot 7 where

the effects of ζ and ν on optimal allocation are significant while Fig. 14 shows examples from ad slot 11 where their effects are not significant.

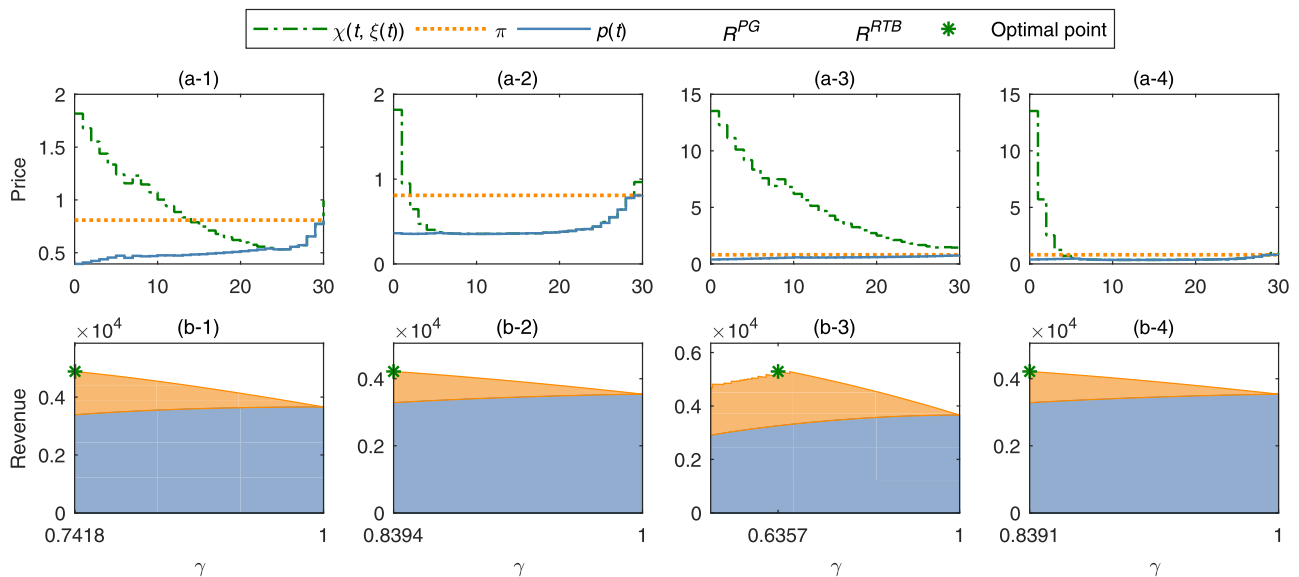


Fig. 13. Examples for ad slot 7: (a) optimal pricing; and (b) allocation of guaranteed contracts. Parameters are set differently in the subplots: (1) $\zeta = 10, \nu = 0.1$; (2) $\zeta = 10, \nu = 0.9$; (3) $\zeta = 90, \nu = 0.1$; (4) $\zeta = 90, \nu = 0.9$.

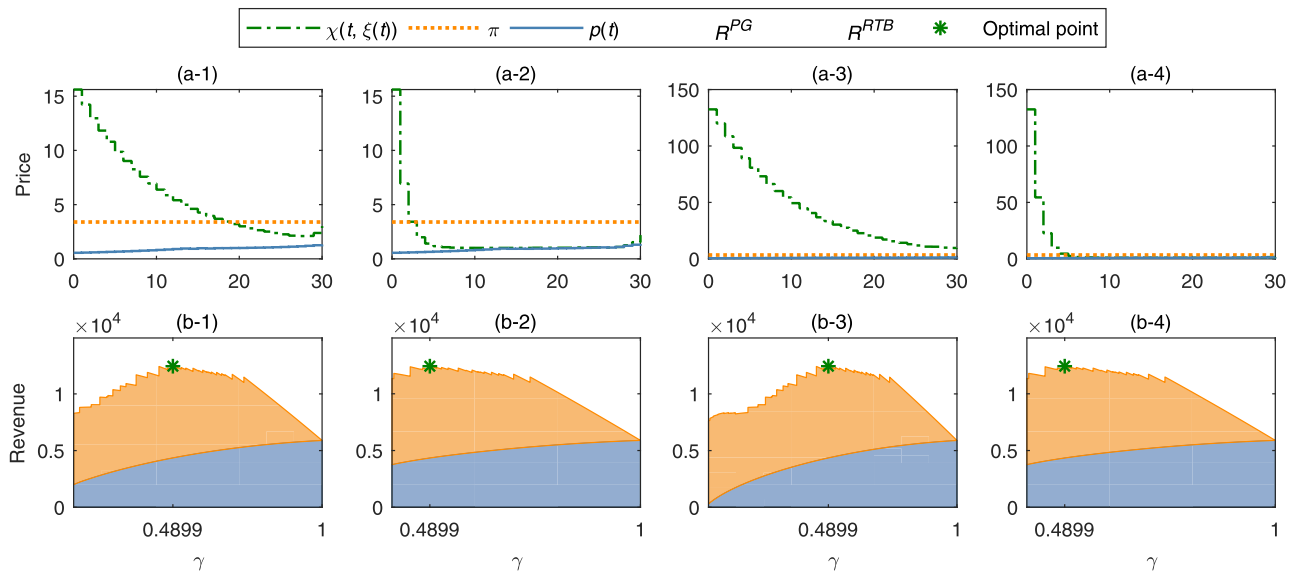


Fig. 14. Examples for ad slot 11: (a) optimal pricing; and (b) allocation of guaranteed contracts. Parameters are set differently in the subplots: (1) $\zeta = 10, \nu = 0.1$; (2) $\zeta = 10, \nu = 0.9$; (3) $\zeta = 90, \nu = 0.1$; (4) $\zeta = 90, \nu = 0.9$.

References

- Abraham, I., Athey, S., Babaioff, M., & Grubb, M. (2013). Peaches, lemons, and cookies: designing auction markets with dispersed information. In *Proceedings of the 14th ACM conference on electronic commerce* (pp. 8–9).
- Anjos, M., Cheng, R., & Currie, C. (2004). Maximizing revenue in the airline industry under one-way pricing. *Journal of the Operational Research Society*, 55(5), 535–541.
- Anjos, M., Cheng, R., & Currie, C. (2005). Optimal pricing policies for perishable products. *European Journal of Operational Research*, 166(1), 246–254.
- Aviv, Y., & Pazgal, A. (2008). Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing & Service Operations Management*, 10(3), 337–352.
- Babaioff, M., Hartline, J., & Kleinberg, R. (2009). Selling ad campaigns: Online algorithms with cancellations. In *Proceedings of the 10th ACM conference on electronic commerce*, (pp. 61–70).
- Balseiro, S., Feldman, J., Mirrokni, V., & Muthukrishnan, M. (2014). Yield optimization of display advertising with ad exchange. *Management Science*, 60(12), 2886–2907.
- Ben-Zwi, O., Henzinger, M., & Loitzenbauer, V. (2015). Ad exchange: Envy-free auctions with mediators. In *Proceedings of the 11th conference web and internet economics*, (pp. 104–117).
- Bharadwaj, V., Chen, P., Ma, W., Nagarajan, C., Tomlin, J., Vassilvitskii, S., et al., (2012). Shale: an efficient algorithm for allocation of guaranteed display advertising. In *Proceedings of the 18th ACM SIGKDD conference on knowledge discovery and data mining*, (pp. 2278–2284).
- Bharadwaj, V., Ma, W., Schwarz, M., Shanmugasundaram, J., Vee, E., Xie, J., et al., (2010). Pricing guaranteed contracts in online display advertising. *Proceedings of the 19th ACM international conference on information and knowledge management*, (pp. 399–408).
- Bishop, C. (2006). *Pattern recognition and machine learning*. Springer.
- Bitran, G., & Caldentey, R. (2003). An overview of pricing models for revenue management. *Manufacturing & Service Operations Management*, 5(3), 179–267.
- Boyle, P. (1986). Option valuation using a three-jump process. *International Options Journal*, 3, 7–12.
- Caldentey, R., & Vulcano, G. (2007). Online auction and list price revenue management. *Management Science*, 53(5), 795–813.
- Chen, B. (2016). Risk-aware dynamic reserve prices of programmatic guarantee in display advertising. In *Proceedings of the 16th IEEE international conference on data mining workshops*, (pp. 511–518).
- Chen, B., & Kankanhalli, M. (2019). Pricing average price advertising options when underlying spot market prices are discontinuous. *IEEE Transactions on Knowledge and Data Engineering*, 31(9), 1765–1778.
- Chen, B., & Wang, J. (2015). A lattice framework for pricing display advertisement options with the stochastic volatility underlying model. *Electronic Commerce Research and Applications*, 14, 465–479.
- Chen, B., Wang, J., Cox, I., & Kankanhalli, M. (2015). Multi-keyword multi-click advertisement option contracts for sponsored search. *ACM Transactions on Intelligent Systems and Technology*, 7(5), <https://boweichen.github.io/publications/TIST2015.pdf>.
- Chen, B., Yuan, S., & Wang, J. (2014). A dynamic pricing model for unifying programmatic guarantee and real-time bidding in display advertising. In *Proceedings of the 8th international workshop on data mining for online advertising* (pp. 1–9).
- Chen, Y. (2017). Optimal dynamic auctions for display advertising. *Operations Research*, 65(4), 897–913.
- Cleveland, W. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74(368), 829–836.
- Constantin, F., Feldman, J., Muthukrishnan, M., & Pál, M. (2009). An online mechanism for ad reservations with cancellations. In *Proceedings of the 20th annual ACM-SIAM symposium on discrete algorithms* (pp. 1265–1274).
- Constantinides, G., & Malliaris, A. (2001). Options markets. In *Equity options markets: Foundations and pricing*: I. Edward Elgar Publishing.
- Cox, J., Ross, S., & Rubinstein, M. (1979). Option pricing: a simplified approach. *Journal of Financial Economics*, 7, 229–263.
- DoubleClick (2005). The decade in online advertising 1994–2004. White Paper, <https://www.slideshare.net/louishatzis/double-click-042005thdecadeinonlineadvertising>.
- Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1), 242–259.
- eMarketer (2013). Real-time bidding poised to make up quarter of all display spending. <http://goo.gl/04Ykx>.
- Feldman, J., Korula, N., Mirrokni, V., Muthukrishnan, M., & Pál, M. (2009). Online ad assignment with free disposal. In *Proceedings of the 5th international workshop on internet and network economics* (pp. 374–385).
- Feldman, J., Mirrokni, V., Muthukrishnan, S., & Pál, M. (2010). Auctions with intermediaries. In *Proceedings of the 11th ACM conference on electronic commerce* (pp. 23–32).
- Fridgeirdottir, K., & Najafi-Asadolahi, S. (2018). Cost-per-impression pricing for display advertising. *Operations Research*, 66(3), 597–892.
- Gallego, G., & van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8), 999–1020.
- Gao, L., Shou, B., Chen, Y.-J., & Huang, J. (2016). Combining spot and futures markets: a hybrid market approach to dynamic spectrum access. *Operations Research*, 64(4), 794–821.
- Ghosh, A., McAfee, P., Papineni, K., & Vassilvitskii, S. (2009). Bidding for representative allocations for display advertising. In *Proceedings of the 5th international workshop on internet and network economics* (pp. 14–18).
- Hojjat, A., Turner, J., Cetintas, S., & Yang, J. (2014). Delivering guaranteed display ads under reach and frequency requirements. In *Proceedings of the 28th AAAI conference on artificial intelligence* (pp. 2278–2284).
- Ilfeld, J., & Winer, R. (2002). Generating website traffic. *Journal of Advertising Research*, 42(5), 49–61.
- Kleinberg, J., & Tardos, E. (2005). *Algorithm design*. Addison Wesley.
- Lahaie, S., & McAfee, P. (2011). Efficient ranking in sponsored search. In *Proceedings of the 7th international conference on internet and network economics* (pp. 254–265).
- Lahaie, S., & Pennock, D. (2007). Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of the 8th ACM conference on electronic commerce* (pp. 50–60).
- Lee, S., & Leckenby, J. (1999). Impact of measurement periods on website rankings and traffic estimation: A user-centric approach. *Journal of Current Issues and Research in Advertising*, 21(2), 1–10.
- Mansour, Y., Muthukrishnan, S., & Nisan, N. (2012). Doubleclick ad exchange auction. arXiv: <https://arxiv.org/abs/1204.0535>.
- McGill, J., & van Ryzin, G. (1999). Revenue management: Research overview and prospects. *Transportation Science*, 33(2), 233–256.

- Muthukrishnan, M. (2009). Ad exchanges: research issues. In *Proceedings of the 5th international workshop on internet and network economics* (pp. 1–12).
- Myerson, R. (1981). Optimal auction design. *Mathematics of Operational Research*, 6(1), 58–73.
- Najafi-Asadolahi, S., & Fridgeirsdottir, K. (2014). Cost-per-click pricing for display advertising. *Manufacturing & Service Operations Management*, 16(4), 482–497.
- Narahari, Y. (2014). *Game theory and mechanism design*. World Scientific.
- OpenX (2013). Programmatic + premium: Current practices and future trends. *White Paper*. <https://www.slideshare.net/digiday/openx-whitepaper-programmatic-premium-current-practices-and-future-trends>.
- Ostrovsky, M., & Schwarz, M. (2011). Reserve prices in Internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM conference on electronic commerce* (pp. 59–60).
- Parkes, D. (2007). *Algorithmic game theory*. chapter Online Mechanisms.. Cambridge University Press.
- Radovanovic, A., & Heavlin, W. (2012). Risk-aware revenue maximization in display advertising. In *Proceedings of the 21st international conference on world wide web* (pp. 91–100).
- Roels, G., & Fridgeirsdottir, K. (2009). Dynamic revenue management for online display advertising. *Journal of Revenue & Pricing Management*, 8(5), 452–466.
- Salomatin, K., Liu, T., & Yang, Y. (2012). A unified optimization framework for auction and guaranteed delivery in online advertising. In *Proceedings of the 21st ACM international conference on information and knowledge management* (pp. 2005–2009).
- Sayedi, A. (2018). Real-time bidding in online display advertising. *Marketing Science*, 37(4), 553–568.
- Su, X. (2007). Intertemporal pricing with strategic customer behavior. *Management Science*, 53(5), 726–741.
- Sun, Z., Dawande, M., Janakiraman, G., & Mookerjee, V. (2016). The making of a good impression: information hiding in ad exchanges. *MIS Quarterly*, 40(3), 717–739.
- Talluri, K., & van Ryzin, G. (2005). *The theory and practice of revenue management*. Springer.
- Thompson, D., & Leyton-Brown, K. (2013). Revenue optimization in the generalized second-price auction. In *Proceedings of the 14th ACM conference on electronic commerce* (pp. 837–852).
- Varian, H. (2007). Position auctions. *International Journal of Industrial Organization*, 25(6), 1163–1178.
- Varian, H. (2009). Online ad auctions. *American Economic Review*, 99(2), 430–434.
- Varian, H., & Harris, C. (2014). The VCG auction in theory and practice. *American Economic Review*, 104(5), 442–445.
- Wang, J., & Chen, B. (2012). Selling futures online advertising slots via option contracts. In *Proceedings of the 21st international world wide web conference* (pp. 627–628).
- Wilmott, P. (2006). *Paul Wilmott on quantitative finance* (2nd). John Wiley.
- Yuan, S., Wang, J., Chen, B., Mason, P., & Seljan, S. (2014). An empirical study of reserve price optimisation in real-time bidding. In *Proceedings of the 20th ACM SIGKDD conference on knowledge discovery and data mining* (pp. 1897–1906).
- Yuan, Y., Wang, F., Li, J., & Qin, R. (2014). A survey on real time bidding advertising. In *Proceedings of the IEEE international conference on service operations and logistics, and informatics* (pp. 418–423).
- Zhang, J., Wang, Z., Li, Q., Zhang, J., Lan, Y., Li, Q., et al. (2017). Efficient delivery policy to minimize user traffic consumption in guaranteed advertising. In *Proceedings of the 31st AAAI conference on artificial intelligence* (pp. 253–258).