

DYNAMIC BIDDING STRATEGIES IN SEARCH-BASED ADVERTISING

SAVAS DAYANIK AND MAHMUT PARLAR

ABSTRACT. Search-based advertising allows the advertisers to run special campaigns targeted to different groups of potential consumers at low costs. Google, Yahoo and Microsoft advertising programs allow the advertisers to bid for an ad position on the result page of a user when the user searches for a keyword that the advertiser relates to its products or services. The expected revenue generated by the ad depends on the ad position, and the ad positions of the advertisers are concurrently determined after an instantaneous auction based on the bids of the advertisers. The advertisers are charged only when their ads are clicked by the users. To avoid excessive ad expenditures due to sudden surges in the keyword-search activities, each advertiser reserves a fixed finite daily budget, and the ads are not shown in the remainder of the day when the budget is depleted. Arrival times of keyword-search instances, ad positions, ad selections, and sales generated by the ads are random. Therefore, an advertiser faces a dynamic stochastic total net revenue optimization problem subject to a strict budget constraint. Here we formulate and solve this problem using stochastic dynamic programming. We show that there is always an optimal dynamic bidding policy. We describe an iterative numerical approximation algorithm that uniformly converges to the optimal solution at an exponential rate of the number of iterations. We illustrate the algorithm on numerical examples. Because the curse of dimensionality of the dynamic programming may sometimes disallow the calculations of the optimal bidding policies, we also propose both static and dynamic alternative bidding policies. We numerically compare the performances of optimal and alternative bidding policies by systematically changing each input parameter. The relative percentage total net revenue losses of the alternative bidding policies increases with budget loading, but were never more than 3.5% of maximum expected total net revenue. The best alternative to the optimal bidding policy turned out to be a static greedy bidding policy.

1. INTRODUCTION

The history of the Internet dates back to the 1950s when point-to-point communication was made possible by mainframe computers and terminals. In later years, as a result of the research conducted by the Defense Advanced Research Projects Agency (DARPA) supported by the US Department of Defense, this type of communication was expanded which allowed communication between (or among) computers. The network funded by DARPA was not commercialized as it was used mainly by universities and other research laboratories in the US. The commercial Internet service providers emerged only in the late 1980s and the Internet was fully commercialized in the early 1990s.

As reported at <http://www.internetworldstats.com/emarketing.htm/>, the Internet users in 1990 totalled about 2.6 million, but by 2010 this number reached 1.6 billion which translates to an annual growth

Date: October 3, 2011.

2000 Mathematics Subject Classification. Primary: 90A60, Secondary: 90C39, 49L20.

OR/MS Classification: Primary: Probability, stochastic model applications. Secondary: dynamic programming/optimal control, applications, finite state Markov.

Key words and phrases. searched-based advertising, dynamic bidding policies, stochastic dynamic programming, optimal control.

rate of about 38%. With the advent of the Internet which became a new source of information and entertainment for millions of people, marketing professionals began paying attention to the new medium to invent new methods of advertising. The first recorded online advertising was a “banner ad” for Hotwire and AT&T in 1994 (Özlük [6]) where an advertisement is embedded into a web page so that it can attract traffic to a website by linking to the website of the advertiser. The “affiliates” which host the advertisements earn money every time a web visitor clicks on the ad.

Banner ads are, however, considered an annoyance by many web users, and their use has been eclipsed by another, and more profitable, form of advertising, i.e., search engine advertisements (also known as keyword ads). These ads appear next to unsponsored (organic) search results, and the advertiser is charged only if a user clicks on the advertisement, taking the user to the advertiser’s webpage. It is estimated by Emarketer.com [2] that keyword ads in the U.S. have generated \$12.4 billion of revenue in 2010 with three quarters of the U.S. market belonging to Google. Although other search engines such as Yahoo! and Bing/MSN operate similarly, in this paper we focus on the advertisers’ bids as submitted to Google.

Suppose a web user is interested in traveling to Italy and visit the countryside on a bicycle, so she searches the keyword “Bike Tours Italy.” As soon as the keyword is entered Google holds an instantaneous generalized second price auction among the advertisers bidding on that keyword and lists, (i) sponsored results, and (ii) organic (unsponsored) search results. If the user clicks on a sponsored link she is directed to the advertiser’s website, and Google charges the advertiser for the click. While there is no guarantee of a sale if the user clicks on a link, Google immediately charges the advertiser for each click which could result in very large bills for the advertiser. To avoid this event, the advertisers are allowed to impose a budget so that if the number of clicks during the day exceeds the budget, the ad is no longer displayed for the rest of the day, even if new visitors search for that keyword “Bike Tours Italy.”

Given the uncertainties involved (i.e., the random arrival of the search requests for a keyword, the random revenue generated from each click, and the random position of the sponsored ad), the problem of how much to bid to maximize the expected profit subject to a budget constraint appears to be a challenging one. In this paper we consider a dynamic version of the keyword bid problem and attempt to determine the optimal bid price(s) subject to the budget constraint using stochastic dynamic programming. It is worth noting that a similar problem has been examined by Cholette, et al. [1] who consider a static version of this problem, i.e., they find the constant optimal bid for each day which maximizes the expected profit under some soft budget constraint. They consider different versions of the problem, i.e., the problem (i) with no constraint, (ii) with a budget constraint where the expected cost of clicks must not exceed the available budget, (iii) with a budget constraint where the probability of exceeding the budget does not exceed a fraction, say, 0.10, and (iv) where the “ideal” bid amount is chosen after selecting a point on the efficient frontier of, (i) expected profit, and (ii) probability of exceeding the budget. Our paper generalizes the Cholette, et al. [1] models by allowing the possibility of *dynamically* selecting the bid prices after observing the remaining budget and the remaining time until the end of the period (usually a day) under hard budget constraint (namely, with probability one, the budget is never exceeded).

In the meantime, Fruchter and Dou [3] solve a different problem with dynamic programming: they find a dynamic strategy to optimally allocate a limited budget between a generic Web portal and a specialized Web portal using keyword-activated banner ads. They concluded that, in the long run, an advertiser must always spend more ad money at the specialized portal.

Kitts and Leblanc [5] formulate an integer program to find optimal bids for multiple keywords on a time grid in an open bid system. They discuss in detail how two key functions appearing in the integer program, namely, the unknown number of clicks as a function of the ad position on the result page and the ad position as a function of the bid amount, can be estimated from the historical data. They report that in a live




FIGURE 1. The spiral perspective (left) and contour plots (middle, right) of a typical optimal dynamic bidding policy function found by the methods of our paper, and 100 simulated budget processes starting at some high (middle) and low (right) initial budget values. Optimal dynamic bidding amounts remain constant with high probability if the initial budget is large (middle), but change quickly if it is not (right).

Overture auction their optimal solution generated four times more web traffic than keyword-specific and rule-based management methods did. Used by *Yahoo's Overture* advertising system, the open bid system allows an advertiser to see the competitors' bids before bidding for the same keyword and updates the bids at some specified frequency. In the *Google's Adwords* bidding system that we focus in this paper, however, the bids are sealed (namely, a bidder cannot see others' bids) and are updated only at the times of the related keyword search. Moreover, the optimal bid amounts are not truly dynamical, because they are concurrently calculated at time zero for all future times and therefore do not utilize the information on the remaining budget amounts, which only reveal immediately before the actual future bidding times. The formulation of this interesting work guarantees only that the expected total cost is less than or equal to the budget. This work is one of the examples with a soft budget constraint in the literature.

Rusmevichientong and Williamson [8] develop an adaptive method of selecting the most profitable keywords at the face of unknown keyword click-through-rates. Their method sorts the keywords in descending profit-to-cost ratios and bids in each period for the maximum number of top keywords that the remaining budget can buy. The keyword click-through-rates are updated at the end of each period. The authors assume that the bid amount, ad position, and cost per click for each keyword remain constant over short periods of time. Our numerical results suggest that the optimal dynamic bid amounts, found by the methods of our paper, remain constant if the initial budget is large, but change quickly over time if the initial budget is small; compare how optimal bidding amounts change along 100 simulated budget processes in Figure 1 for large (middle figure) and low (right figure) initial budget values. More in-dept discussion of Figure 1 is given below in this section as we outline our own contributions.

Özlük and Cholette [7] also assume that one bid price is used for each keyword over the entire time period. By solving a constrained nonlinear optimization problem, they find the optimal bidding amounts that maximize the expected total net revenue subject to a soft budget constraint: the expected total cost should not exceed the budget. Later Cholette et al. [1] enhanced the same model by explicitly modeling the stochastic ad positions, which we are adopting here as well, and by considering a probabilistic budget exceedance constraint, which is also a soft budget constraint.

Under soft budget constraints, the *expected* total cost must be less than the available budget or the probability that the total cost exceeds the available budget must be small, but the actual spendings may

occasionally exceed the budget with positive probability. In reality, however, the ad is not displayed in the remainder of the period after the budget is depleted. Therefore, optimal bidding under soft budget constraints will overestimate the true maximum expected total net revenue. In our paper, we solve the optimal bidding problem under a strict budget constraint and calculate the un-inflated maximum expected total net revenue.

The strict budget constraint and stochastic ad positions also force the optimal strategy to *dynamically respond to stochastically fluctuating budget process*: if the budget is underutilized in excess amounts for a while due to slow click traffic, then optimal strategy should raise the bid amounts in order to drive ad positions and click-through rates up. If the budget is overutilized because of frequent clicks, then the optimal strategy should check if the same traffic can be generated with lower bid amounts and lower the bid amounts if necessary. The optimal bidding strategy should therefore depend on the remaining budget and the remaining time. It should balance the trade-off between the maximum traffic (and therefore revenue) that can be generated in the remainder of time and the total cost of doing that. The highly nonlinear relation between bid amount, ad position, and ad clickability make impossible to further speculate on the behavior of optimal bidding strategies. Our stochastic dynamic programming solution first of all proves the existence of an optimal bidding strategy and provides an exponentially fast uniformly converging algorithm to find provably nearly-optimal bidding strategies. Moreover, the numerical illustrations uncover the simple very intuitive spiral look of optimal bidding amounts on the space spanned by the remaining budget and remaining time variables; see Figure 1. The optimal bidding amounts increase along any horizontal line drawn on the contour plot as the line is traced from right to left. Namely, if for a long time no clicks are received and budget does not decrease, then the optimal bid amounts increase. The optimal bidding amounts decrease along any vertical line drawn on the contour plot as the line is traced from top to down. Namely, if excessive number of clicks are received over a very short time period, then the optimal bid amounts decrease. Both observations are consistent with our expectations.

Despite the fact that the stochastic dynamic programming provides optimal solution and valuable insights about the structure of the optimal bidding strategy, the computational time grows with the size of the grid on budget-time space. We therefore propose and then examine two intuitive dynamic bidding strategies that can be calculated fast. The first strategy is the dynamical version of “budget-constrained” (BC) bidding strategy of Cholette et al. [1]. S. Cholette and her colleagues proposed a *static BC bidding strategy* which bids all the time the same amount that initially maximizes the expected total net revenue subject to the constraint that the expected total cost should not exceed the budget. They show that the problem can easily be solved with Lagrangian relaxation. Here we propose the *dynamic BC (DBC) bidding strategy*, which recalculates the BC bid amount after the values of the remaining budget and time in the budget constraint are updated after every keyword search. We alternatively define a *static greedy (G) bidding strategy* that bids all the time the same bid amount that maximize the expected total net revenue subject to strict budget constraint. Finding this amount boils down to solving an unconstrained nonlinear optimization problem on a compact bounded interval. We finally introduce the *dynamic greedy (DG) bidding strategy* that recalculates the greedy bid amount after the remaining budget and remaining time in the budget constraint are updated following every keyword search.

In an extensive numerical study, we compared the performances of optimal dynamic bidding policy, optimal static and dynamic BC bidding policies, and optimal static and dynamic greedy bidding strategies. The dynamic greedy bidding strategy performs slightly better than the dynamic BC bidding strategy in general. Both strategies however perform quite well and achieve expected total net revenues not falling short more than 3.5% of that of the optimal bidding strategy in the numerical experiments. We noticed that the largest revenue losses occur at high budget loadings (fraction of the initial budget spent on bidding). Figure

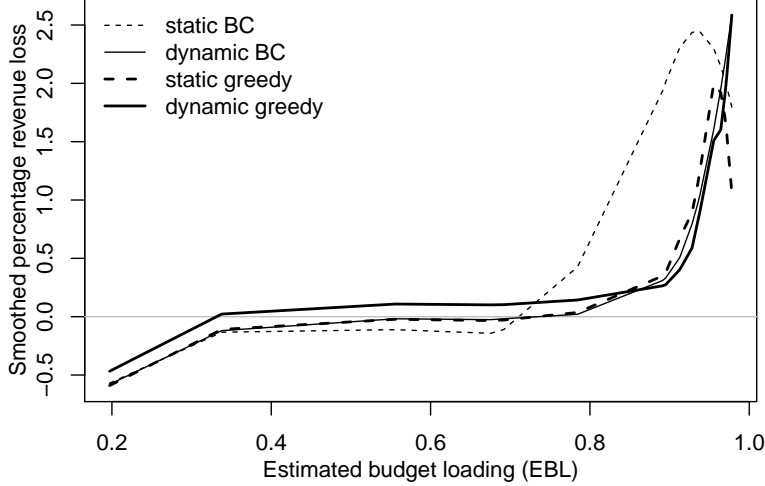


FIGURE 2. Comparison of static and dynamic bidding policies

2 shows the *smoothed* percentage revenue losses of the static/dynamic BC/greedy bidding policies relative to the optimal dynamic bidding policy (see Figure 7 for individual data points). All policies perform close to optimal if the budget loading is less than $0.7 \sim 0.8$. As the budget loading increases to 1, the percentage revenue losses sharply increase. The simple static greedy bidding policy performs as good as the dynamic BC and greedy bidding policies and outperforms the static BC bidding policy. Moreover, unlike for the dynamic BC and greedy policies, the percentage revenue losses for static greedy policy seems to be a bounded function of budget loading. When the optimal dynamic bidding policy is difficult to calculate or has to be frequently recalculated because of changing/reestimated parameters, the static greedy bidding policy looks the best alternative among all other static and dynamic bidding policies studied in this paper.

In Section 2, we formulate the dynamic bidding problem and solve it with stochastic dynamic programming. We show that the dynamic bidding problem always admits an optimal bidding policy. After we define a value function and heuristically derive a dynamic programming equation, we show that the equation has unique solution and verify that the solution coincides with the value function of the dynamic bidding problem. The verification is done with appropriate martingales. In the meantime, we derive decreasing and increasing successive approximations of the value function, and they lead to an iterative numerical solution algorithm that converges uniformly and at an exponential rate. Section 3 introduces simple and fast static/dynamic budget-constrained/greedy bidding policies as alternatives to optimal bidding policy for those cases when the dynamic programming's curse of dimensionality disallows the calculation of optimal bidding policy. In Section 4, the numerical solution algorithm is described and illustrated on several examples. The maximum total expected net revenues for optimal and BC/greedy policies are systematically compared as each of the parameters B, a, λ, m, μ is changed one at a time over a five-point grid. The relative percentage losses of static/dynamic BC/greedy policies are compared as the estimated budget loading changes. Finally, the lemmas we used in Section 2 are given in the appendix.

2. PROBLEM FORMULATION AND SOLUTION

We formulate the optimal bidding problem as it is faced by an advertiser who works with a web search company similar to Google. When a keyword is searched by a search-engine user, a generalized second price auction immediately takes place among all advertisers whose lists contain the keyword. Each bid is sealed and is not seen by the others. Every bid is then multiplied by some quality score of the bidder, which is

proportional to the bidder's ability to consistently generate high internet traffic. Here we assume that all of the advertisers take all necessary steps to get top quality scores (e.g., relevant keywords are carefully chosen, ads are attractive, landing web pages are well designed) so that after multiplication by the quality scores the ordering of the bids does not change. Google [4] provides detailed information about how the quality score is calculated and how it can be improved; see also Szetela and Kerschbaum [9, Chapter 9].

The bids of all advertisers are then sorted in descending order, and their messages are placed on the result page according to the same order. Therefore, the ad position of each ad is a random variable and determines the probability that the user will click that ad. The closer to the top of the result page the ad is, the more likely is the ad to be clicked by the user. The advertiser is not charged unless the user clicks the ad.

According to the generalized second price auction, each bidder pays the largest bid received for the same keyword less than or equal to her own bid. However, the difference between those bids is often negligible, and therefore, we assume that each bidder pays her bid amount if the user clicks her ad. Finally, we assume that the revenue realized after each click is a random variable independent of the bid amount, ad position, number and times of keyword search events.

Let $X(t)$ be the number of keywords in our advertiser's list typed in $(0, t]$ for every $t \geq 0$. We assume that the users from all over the world search for the keywords independently of each other. Because the Internet services are available throughout the day and one part of the world is always awake even if the other parts are in rest, we also assume that the keyword search activity resumes on average at the same pace on any typical day. Therefore, we assume that $X = \{X(t); t \geq 0\}$ follows a homogeneous Poisson process with some constant rate $\lambda > 0$ and with arrival times T_1, T_2, \dots . The advertiser starts with some initial budget B_0 and bids b_n for an ad position on the result page that will be generated by the n th keyword search. Let B_n be the remaining budget before the n th keyword search.

Let L_n be the ad position on the result page and Z_n be 1/0 if the ad is clicked/not clicked after the n th keyword search. We denote the top and bottom of a result page by 0 and 1, respectively, and identify the entire page with the interval $[0, 1]$. Hence, L_n is a random variable taking values in $[0, 1]$. Then we have

$$B_n = B_{n-1} - Z_{n-1}b_{n-1}, \quad n \geq 1,$$

and we assume that

$$\begin{aligned} \mathbb{P}\{Z_n = 1 \mid L_n = \ell\} &= (p_0 - p_1)(1 - \ell)^m + p_1, & \ell \in [0, 1], n \geq 0, \\ \mathbb{P}\{L_n \in d\ell \mid b_n = b\} &= \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \ell^{a-1} (1 - \ell)^{b-1} d\ell, & \ell \in [0, 1], n \geq 0 \end{aligned}$$

for some real numbers $0 \leq p_1 \leq p_0 \leq 1$, $m \geq 0$, and $a > 0$. The probability that an ad is clicked by the user decreases with the ad's position on the result page. This probability takes its maximum value p_0 if the ad is placed at the top and its minimum value p_1 if it is placed at the bottom of the page. The parameter a may be thought as the average bid amount of all other competing advertisers bidding on the same keyword. Because in a sealed bidding scheme an advertiser cannot see the other bids, this representation of the competitors' average bidding behavior is reasonable. The ad position is a Beta distributed random variable with parameters a and b (our advertiser's bid amount). As b increases (respectively, decreases), the ad is more likely to be placed closer to the top (respectively, bottom) of the result page. Figure 3 shows some examples of the click probability and ad position's probability density function as, respectively, ad position and bid amount change.

Suppose finally that the users potentially generate revenues W_1, W_2, \dots , which are i.i.d. random variables with common finite mean μ independent of Z_1, Z_2, \dots . A potential revenue W_i realizes if and only if the i th user clicks the ad; namely, if and only if $Z_i = 1$, for every $i \geq 0$.

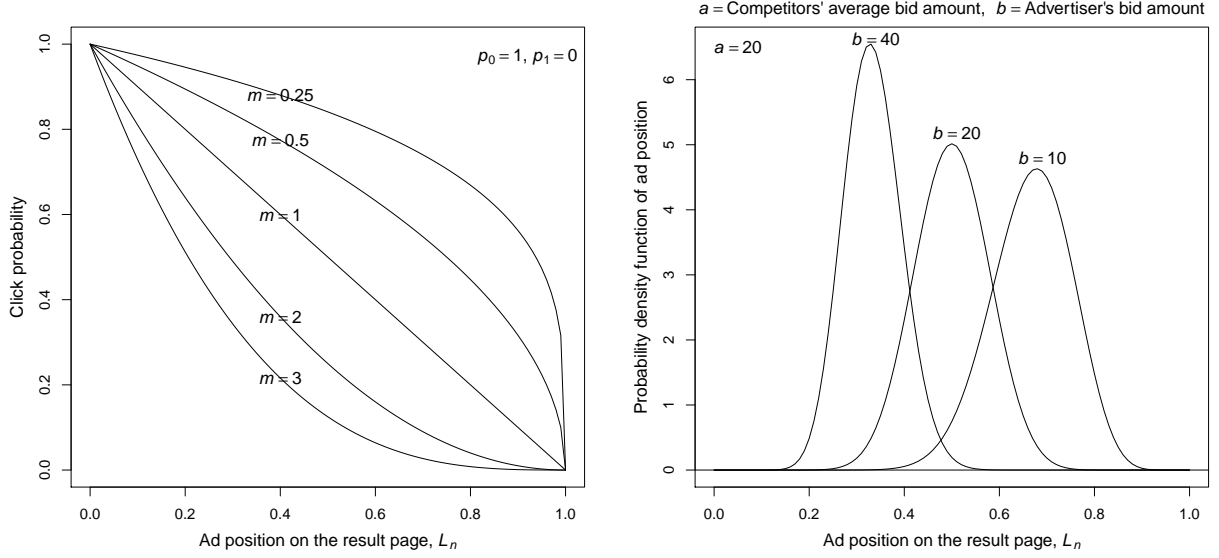


FIGURE 3. Click probability as a function of ad position and the probability density function of ad position on the result page as the advertiser's bid amount changes

2.1. The dynamic programming formulation. Assume that the first keyword search takes place at time 0. We make this assumption in order to be able to evoke the optimality principle and will later relax it. Let $V(B, T)$ be the maximum expected total net revenue obtainable with budget B to be used in the remaining T units of time. Then

$$(1) \quad V(B, T) = \sup_{(b_i)_{i \geq 0} \in \mathcal{D}(B, T)} \mathbb{E} \left[\sum_{i=0}^{\infty} (W_i - b_i) Z_i 1_{\{T_i \leq T\}} \right],$$

where the supremum is taken over the collection of *admissible* bidding strategies,

$$\mathcal{D}(B, T) = \left\{ (b_i)_{i \geq 0}; b_i \geq 0 \text{ for every } i \geq 0, \text{ and } \sum_{i=0}^{\infty} b_i Z_i 1_{\{T_i \leq T\}} \leq B \right\}.$$

Suppose that we are interested in bidding problem for every $0 \leq B \leq B_{\max}$ and $0 \leq T \leq T_{\max}$ for any arbitrary but fixed and finite $B_{\max} > 0$ and $T_{\max} > 0$. The problem is to find $V(B, T)$ and an optimal bidding strategy, if there is one, for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$. The principle of optimality of the dynamic programming suggests that

$$(2) \quad V(B, T) = \sup_{0 \leq b \leq B} \mathbb{E} \left[(W_0 - b) Z_0 + V(B - Z_0 b, T - T_1) 1_{\{T_1 \leq T\}} \right], \quad B, T \geq 0$$

with the boundary conditions $V(0, T) = 0$ for every $T \geq 0$ and

$$V(B, 0) = \max_{0 \leq b \leq B} \mathbb{E}[(W_0 - b) Z_0] = \max_{0 \leq b \leq B} (\mu - b) G(b) \quad \text{for every } B \geq 0,$$

where the conditional click probability, given the bid amount, is denoted by

$$\begin{aligned} G(b) &:= \mathbb{P}\{Z_0 = 1 \mid b_0 = b\} = \int \mathbb{P}\{Z_0 = 1 \mid L_0 = \ell\} \mathbb{P}\{L_0 \in d\ell \mid b_0 = b\} \\ &= \int_0^1 [(p_0 - p_1)(1 - \ell)^m + p_1] \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \ell^{a-1} (1 - \ell)^{b-1} d\ell = (p_0 - p_1) \frac{\Gamma(a+b)\Gamma(b+m)}{\Gamma(a+b+m)\Gamma(b)} + p_1 \\ &\approx (p_0 - p_1) \sqrt{\frac{(a+b+m)b}{(a+b)(b+m)}} \left(\frac{a+b}{a+b+m} \right)^{a+b} \left(\frac{b+m}{b} \right)^b \left(\frac{b+m}{a+b+m} \right)^m + p_1, \end{aligned}$$

which follows from the Stirling's approximation $\Gamma(z) \approx \sqrt{2\pi/z} (z/e)^z$, $z > 0$ of the Gamma function, where $f(x) \approx g(x)$ in the sense that $\lim_{x \uparrow \infty} f(x)/g(x) = 1$. If $a+b+m$ is large, then the numerator and denominator of the exact expression for $G(b)$ can be extremely large. The approximation limits the potential numerical problems because the quotients inside power functions are either less than or close to one as long as b is small or $m \ll b$.

Because $(b_i)_{i \geq 0}$ with $b_i = 0$, $i \geq 0$ is in $\mathcal{D}(B, T)$, and $G(0) = 0$ implies that under bidding strategy $(b_i)_{i \geq 0}$ we have $Z_0 = Z_1 = \dots = 0$ and obtain zero net revenue with probability one, we have $V(B, T) \geq 0$ for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$. On the other hand,

$$\mathbb{E} \left[\sum_{i=0}^{\infty} (W_i - b_i) Z_i 1_{\{T_i \leq T\}} \right] \leq \mathbb{E} \left[\sum_{i=0}^{\infty} W_i 1_{\{T_i \leq T\}} \right] = \mathbb{E} W_0 [1 + \mathbb{E} X(T)] = \mu(1 + \lambda T).$$

Hence, the value function in (1) is finite, and

$$(3) \quad 0 \leq V(B, T) \leq \mu(1 + \lambda T) \quad \text{for every} \quad (B, T) \in [0, B_{\max}] \times [0, T_{\max}].$$

We shall show that i) the problem (1) always admits an optimal bidding strategy, and ii) the value function $V(\cdot, \cdot)$ and an optimal bidding strategy attaining the supremum in (1) can be calculated numerically with successive approximations.

2.2. A dynamic programming operator and successive approximations. Let $\mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$ be the collection of all continuous functions from $[0, B_{\max}] \times [0, T_{\max}]$ to \mathbb{R} , equipped with sup-norm $\|w\| = \sup_{(B, T) \in [0, B_{\max}] \times [0, T_{\max}]} |w(B, T)|$ for every $w \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$. Let M be an operator acting on \mathcal{C} defined by

$$(4) \quad (Mw)(B, T) := \sup_{0 \leq b \leq B} \mathbb{E} \left[(W_0 - b) Z_0 + w(B - Z_0 b, T - T_1) 1_{\{T_1 \leq T\}} \right]$$

for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$. Since the form of the righthand side is suggested by the dynamic programming optimality principle, M is often called a dynamic programming operator. Because $\mathbb{E}[(W_0 - b) Z_0 + w(B - Z_0 b, T - T_1) 1_{\{T_1 \leq T\}}]$ equals

$$\begin{aligned} (5) \quad (M_1 w)(b, B, T) &:= \mathbb{E}[(W_0 - (b \wedge B) Z_0 + w(B - Z_0(b \wedge B), T - T_1) 1_{\{T_1 \leq T\}})] \\ &= [\mu - (b \wedge B)] G(b \wedge B) + \int_0^T \lambda e^{-\lambda t} \left[w((B - b)^+, T - t) G(b \wedge B) + w(B, T - t) (1 - G(b \wedge B)) \right] dt \\ &= \int_0^T \lambda e^{-\lambda t} w(B, T - t) dt + G(b \wedge B) \left\{ \mu - (b \wedge B) + \int_0^T \lambda e^{-\lambda t} \left[w((B - b)^+, T - t) - w(B, T - t) \right] dt \right\}, \end{aligned}$$

which is defined for all $0 \leq b \leq B_{\max}$, we can write (4) more explicitly as

$$(6) \quad (Mw)(B, T) = \sup_{0 \leq b \leq B_{\max}} (M_1w)(b, B, T) = \sup_{0 \leq b \leq B} (M_1w)(b, B, T) = \int_0^T \lambda e^{-\lambda t} w(B, T-t) dt \\ + \sup_{0 \leq b \leq B} G(b) \left\{ \mu - b + \int_0^T \lambda e^{-\lambda t} [w(B-b, T-t) - w(B, T-t)] dt \right\}.$$

We shall show that M is a contraction mapping, and that the value function $V(\cdot, \cdot)$ in (1) is the unique fixed point of operator M , and that $V(\cdot, \cdot)$ can be calculated as the limit of the increasing $(\underline{V})_{n \geq 0}$ or decreasing $(\bar{V})_{n \geq 0}$ successive approximations, respectively defined by

$$(7) \quad \underline{V}_0(B, T) \equiv 0 \quad \text{and} \quad \underline{V}_n(B, T) := (M\underline{V}_{n-1})(B, T), \quad (B, T) \in [0, B_{\max}] \times [0, T_{\max}], \\ \bar{V}_0(B, T) = \mu(1 + \lambda T) \quad \text{and} \quad \bar{V}_n(B, T) := (M\bar{V}_{n-1})(B, T), \quad (B, T) \in [0, B_{\max}] \times [0, T_{\max}].$$

Below Proposition 1 and Theorem 1 identify some key properties of operator M and successive approximations $(\underline{V})_{n \geq 0}$, $(\bar{V})_{n \geq 0}$ that will be needed to show the existence of an optimal bidding policy and a uniformly and exponentially convergent numerical algorithm to calculate both the value function and nearly-optimal bidding policies.

Proposition 1. *Let w, w_1, w_2 be in $\mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$.*

- (i) $M_1w \in \mathcal{C}([0, B_{\max}]^2 \times [0, T_{\max}])$ and $Mw \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$.
- (ii) All three suprema in (6) are attained.
- (iii) If $w_1 \leq w_2$, then $Mw_1 \leq Mw_2$. If $w(B, T) \leq \mu(1 + \lambda T)$ for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$, then $(Mw)(B, T) \leq \mu(1 + \lambda T)$ for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$.
- (iv) M is a contraction mapping on $\mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$ and $\|Mw_1 - Mw_2\| \leq (1 - e^{-\lambda T_{\max}})\|w_1 - w_2\|$. If M has a fixed point, then it must be unique.

Proof. (i) The continuity of M_1w and Mw is proved with Lemma 4 in the appendix. (ii) Because $b \mapsto (M_1w)(b, B, T)$ is continuous on compact $[0, B_{\max}]$, all three suprema in (6) are attained. (iii) The monotonicity of operator M in w immediately follows from its definition in (4). Suppose now that $w(B, T) \leq \mu(1 + \lambda T)$ for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$. By the memorylessness property of exponentially distributed inter-arrival times of X ,

$$\mathbb{E}[(T - T_1)1_{\{T_1 \leq T\}}] = \mathbb{E}[(T - T_1)(1 - 1_{\{T_1 > T\}})] = \mathbb{E}[T - T_1] - \mathbb{E}[(T - T_1)1_{\{T_1 > T\}}] \\ = T - \mathbb{E}T_1 + \mathbb{E}[T_1 - T \mid T_1 > T] \mathbb{P}\{T_1 > T\} = T - \mathbb{E}T_1 + \mathbb{E}T_1 \mathbb{P}\{T_1 > T\} = T - (1 - e^{-\lambda T})/\lambda,$$

and

$$(Mw)(B, T) := \sup_{0 \leq b \leq B} \mathbb{E}[(W_0 - b)Z_0 + w(B - Z_0b, T - T_1)1_{\{T_1 \leq T\}}] \leq \mathbb{E}[W_0 + \mu(1 + \lambda(T - T_1))1_{\{T_1 \leq T\}}] \\ = \mu + \mu \mathbb{P}\{T_1 \leq T\} + \mu \lambda \mathbb{E}[(T - T_1)1_{\{T_1 \leq T\}}] = \mu + \mu(1 - e^{-\lambda T}) + \mu \lambda T - \mu(1 - e^{-\lambda T}) = \mu(1 + \lambda T)$$

for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$. (iv) To show that M is a contraction mapping, take any $w_1, w_2 \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$ and any $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$. By the first part, there are $b_1, b_2 \in [0, B]$ at which the suprema in (6) are attained for $w = w_1$ and $w = w_2$, respectively. Then

$$(Mw_1)(B, T) - (Mw_2)(B, T) \leq (M_1w_1)(b_1, B, T) - (M_1w_2)(b_1, B, T) \\ = \mathbb{E}[(W_0 - b_1)Z_0 + w_1(B - Z_0b_1, T - T_1)1_{\{T_1 \leq T\}}] - \mathbb{E}[(W_0 - b_1)Z_0 + w_2(B - Z_0b_1, T - T_1)1_{\{T_1 \leq T\}}] \\ \leq \mathbb{E}[|w_1(B - Z_0b_1, T - T_1) - w_2(B - Z_0b_1, T - T_1)|1_{\{T_1 \leq T\}}] \leq (1 - e^{-\lambda T_{\max}})\|w_1 - w_2\|.$$

If we interchange the roles of w_1 and w_2 and replace b_1 with b_2 above, then we obtain $(Mw_2)(B, T) - (Mw_1)(B, T) \leq (1 - e^{-\lambda T_{\max}})\|w_2 - w_1\|$. Hence, $|(Mw_1)(B, T) - (Mw_2)(B, T)| \leq (1 - e^{-\lambda T_{\max}})\|w_1 - w_2\|$ for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$, which implies that $\|Mw_1 - Mw_2\| \leq (1 - e^{-\lambda T_{\max}})\|w_1 - w_2\|$ and M is a contraction mapping on $\mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$. Finally, suppose that w_1, w_2 are two fixed points of M ; namely, $w_1 = Mw_1$ and $w_2 = Mw_2$. Then $\|w_1 - w_2\| = \|Mw_1 - Mw_2\| \leq (1 - e^{-\lambda T_{\max}})\|w_1 - w_2\|$ implies that $\|w_1 - w_2\| = 0$ and $w_1 = w_2$. \square

The relation (1) suggested by the optimality principle implies that the best candidate for the value function (1) is a fixed point of operator M . The plan is to find the fixed points of operator M and then verify if any of them coincides with the value function.

Theorem 1 (iv) below proves that M has unique continuous fixed point. We leave to Section 2.3 the verification that the only fixed point of M is indeed the value function of (1). In applications, however, it is also important to know how to calculate the fixed point of M and hence the value function (1). The remainder parts of Theorem 1 state that the fixed point is the common limit point of the decreasing and increasing successive approximations in (7). Moreover, for any arbitrarily small choice of a positive tolerance number, both sequences can be terminated finitely and two approximations obtained in this way are pointwise different from the fixed point not more than the preset tolerance number.

Theorem 1. *Let $(\underline{V}_n)_{n \geq 0}$ and $(\bar{V}_n)_{n \geq 0}$ be the sequences of successive approximations defined by (7).*

- (i) $\underline{V}_n, \bar{V}_n \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$ for every $n \geq 0$.
- (ii) $0 = \underline{V}_0(B, T) \leq \underline{V}_1(B, T) \leq \dots \leq \bar{V}_1(B, T) \leq \bar{V}_0(B, T) = \mu(1 + \lambda T)$, $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$.
- (iii) *The limits $\underline{V}_\infty(B, T) := \uparrow \lim_{n \rightarrow \infty} \underline{V}_n(B, T)$ and $\bar{V}_\infty(B, T) := \downarrow \lim_{n \rightarrow \infty} \bar{V}_n(B, T)$ exist for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$, and*

$$\begin{aligned} \|\underline{V}_\infty - \underline{V}_n\| &\leq \mu(1 + \lambda T_{\max})e^{\lambda T_{\max}}(1 - e^{-\lambda T_{\max}})^n \quad \text{for every } n \geq 0, \\ \|\bar{V}_\infty - \bar{V}_n\| &\leq \mu(1 + \lambda T_{\max})e^{\lambda T_{\max}}(1 - e^{-\lambda T_{\max}})^n \quad \text{for every } n \geq 0. \end{aligned}$$

- (iv) $\underline{V}_\infty = \bar{V}_\infty \equiv V_\infty \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$ is the unique fixed point of M ; namely, $V_\infty = MV_\infty$.
- (v) For every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$, the suprema

$$\begin{aligned} \underline{V}_n(B, T) &= (M\underline{V}_{n-1})(B, T) = \sup_{0 \leq b \leq B} (M_1 \underline{V}_{n-1})(b, B, T), \quad n \geq 1, \\ \bar{V}_n(B, T) &= (M\bar{V}_{n-1})(B, T) = \sup_{0 \leq b \leq B} (M_1 \bar{V}_{n-1})(b, B, T), \quad n \geq 1, \\ V_\infty(B, T) &= (MV_\infty)(B, T) = \sup_{0 \leq b \leq B} (M_1 V_\infty)(b, B, T) \end{aligned}$$

are attained; namely, there are some $\underline{b}_n(B, T), \bar{b}_n(B, T) \in [0, B]$, $n \geq 1$ and $b_\infty(B, T) \in [0, B]$ such that $\underline{V}_n(B, T) = (M_1 \underline{V}_{n-1})(\underline{b}_n(B, T), B, T)$, $n \geq 1$ and $\bar{V}_n(B, T) = (M_1 \bar{V}_{n-1})(\bar{b}_n(B, T), B, T)$, $n \geq 1$ and $V_\infty(B, T) = (M_1 V_\infty)(b_\infty(B, T), B, T)$.

Proof. Because $\underline{V}_0, \bar{V}_0 \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$ and $\underline{V}_0(B, T), \bar{V}_0(B, T) \leq \mu(1 + \lambda T)$ for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$, Proposition 1 (i & iii) and successive applications of M imply both (i) and (ii). Therefore, the pointwise limits $\underline{V}_\infty = \uparrow \lim_{n \rightarrow \infty} \underline{V}_n$ and $\bar{V}_\infty = \downarrow \lim_{n \rightarrow \infty} \bar{V}_n$ obviously exist. We shall prove (iii)-(v) for $(\underline{V}_n)_{n \geq 0}$ and \underline{V}_∞ ; the proofs are identical for $(\bar{V}_n)_{n \geq 0}$ and \bar{V}_∞ . Because $\|\underline{V}_{k+1} - \underline{V}_k\| = \|M\underline{V}_k - M\underline{V}_{k-1}\| \leq (1 - e^{-\lambda T_{\max}})\|\underline{V}_k - \underline{V}_{k-1}\| \leq \dots \leq (1 - e^{-\lambda T_{\max}})^k \|\underline{V}_1\|$ by Proposition 1 (iv), for every $m > n$ we have

$$\|\underline{V}_m - \underline{V}_n\| \leq \sum_{k=n}^{m-1} \|\underline{V}_{k+1} - \underline{V}_k\| \leq \sum_{k=n}^{m-1} (1 - e^{-\lambda T_{\max}})^k \|\underline{V}_1\|$$

Passing to limit as $m \rightarrow \infty$, we get that $\|\underline{V}_\infty - \underline{V}_n\| \leq \sum_{k=n}^{\infty} (1 - e^{-\lambda T_{\max}})^k \|\underline{V}_1\|$, which together with part (ii) leads to (iii). Because the inequality in part (iii) implies $\|\underline{V}_\infty - \underline{V}_n\| \xrightarrow{n \rightarrow \infty} 0$, the continuous functions \underline{V}_n , $n \geq 0$ converge to \underline{V}_∞ uniformly. Therefore, $\underline{V}_\infty \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$ and $\underline{V}_\infty = \lim_{n \rightarrow \infty} \underline{V}_{n+1} = \lim_{n \rightarrow \infty} M \underline{V}_n = M(\lim_{n \rightarrow \infty} \underline{V}_n) = M \underline{V}_\infty$. Thus, \underline{V}_∞ is continuous and is a fixed point of M . (v) follows from the continuity of $(\underline{V}_n)_{n \geq 0}$, \underline{V}_∞ and Proposition 1 (ii). Finally, because \bar{V}_∞ is also a fixed point of M , which is unique by Proposition 1 (iv), we must have $\underline{V}_\infty = \bar{V}_\infty$, which is denoted in the remainder by V_∞ . \square

2.3. The verification of the dynamic programming solution and the existence of an optimal bidding policy. We will now show that the function V_∞ , which is the unique fixed point of M in $\mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$ and the common pointwise uniform limit of the increasing successive approximations $(\underline{V}_n)_{n \geq 0}$ and the decreasing successive approximations $(\bar{V}_n)_{n \geq 0}$, coincides with the value function V in (1). Theorem 3 is the main result and is a verification theorem based on martingales, which we now introduce:

Let $\mathbf{b} = (b_n)_{n \geq 0} \in \mathcal{D}(B, T)$ be any admissible bidding strategy and $(\mathcal{F}_n^{\mathbf{b}})_{n \geq 0}$ be the corresponding observation filtration, where

$$\mathcal{F}_n^{\mathbf{b}} = \sigma\{B_0, T_0, \dots, B_n, T_n\}, \quad n = 0, 1, \dots$$

Then $B_n, b_n \in \mathcal{F}_n^{\mathbf{b}}$, $Z_n = B_{n+1} - B_n \in \mathcal{F}_{n+1}^{\mathbf{b}}$, and W_n is independent of $(\mathcal{F}_i^{\mathbf{b}})_{i \geq 0}$ for every $n \geq 0$. Let

$$Y_0^{\mathbf{b}} := V_\infty(B_0, T) \quad \text{and} \quad Y_n^{\mathbf{b}} := \sum_{k=0}^{n-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} + V_\infty(B_n, T - T_n) 1_{\{T_n \leq T\}}, \quad n \geq 0.$$

The process $(Y_n^{\mathbf{b}})_{n \geq 0}$ is adapted to the filtration $(\mathcal{F}_n^{\mathbf{b}})_{n \geq 0}$; namely, $Y_n^{\mathbf{b}} \in \mathcal{F}_n^{\mathbf{b}}$ for every $n \geq 0$. It is the dynamic representation of how bright the past and future look under bidding policy \mathbf{b} after the $(n-1)$ st keyword-search instance. Theorem 2 states that, unless the advertiser follows the bidding policy $\mathbf{b}^* \in \mathcal{D}(B, T)$, the total net revenue may only slide down over time from its maximum possible value. This strongly hints that \mathbf{b}^* is an optimal bidding strategy.

Theorem 2. *For any bidding strategy $\mathbf{b} = (b_n)_{n \geq 0} \in \mathcal{D}(B, T)$, the process $(Y_n^{\mathbf{b}}, \mathcal{F}_n^{\mathbf{b}})_{n \geq 0}$ is a supermartingale. If $\mathbf{b}^* = (b_n^*)_{n \geq 0}$ is the admissible bidding strategy defined by $b_n^* := b_\infty(B_n, T - T_n) 1_{\{T_n \leq T\}}$ for every $n \geq 0$, where $b_\infty(B, T) := \arg \max_{0 \leq b \leq B} (M_1 V_\infty)(b, B, T)$ as in Theorem 1, then $(Y_n^{\mathbf{b}^*}, \mathcal{F}_n^{\mathbf{b}^*})_{n \geq 0}$ is a martingale.*

Proof. For every $n \geq 0$, $Y_n^{\mathbf{b}} \in \mathcal{F}_n^{\mathbf{b}}$ is integrable, and $Y_{n+1}^{\mathbf{b}} - Y_n^{\mathbf{b}}$ equals

$$\begin{aligned} & (\mu - b_n) Z_n 1_{\{T_n \leq T\}} + V_\infty(B_{n+1}, T - T_{n+1}) 1_{\{T_{n+1} \leq T\}} - V_\infty(B_n, T - T_n) 1_{\{T_n \leq T\}} \\ & + V_\infty(B_n - b_n Z_n, T - T_n - (T_{n+1} - T_n)) 1_{\{T_n \leq T\}} 1_{\{T_{n+1} - T_n \leq T - T_n\}} - V_\infty(B_n, T - T_n) 1_{\{T_n \leq T\}}. \end{aligned}$$

Because the random variables $B_n, b_n, T - T_n$ are $\mathcal{F}_n^{\mathbf{b}}$ -measurable, and $\mathbb{E}[Z_n | \mathcal{F}_n^{\mathbf{b}}] = \mathbb{P}\{Z_n = 1 | b_n\} = G(b_n)$, and $T_{n+1} - T_n$ has exponential distribution with rate λ independently of $\mathcal{F}_n^{\mathbf{b}}$, we have

$$\begin{aligned} (8) \quad \mathbb{E}[Y_{n+1}^{\mathbf{b}} - Y_n^{\mathbf{b}} | \mathcal{F}_n^{\mathbf{b}}] &= 1_{\{T_n \leq T\}} \left\{ (\mu - b_n) G(b_n) + \int_0^{T-T_n} \lambda e^{-\lambda t} [V_\infty(B_n - b_n, T - T_n - t) G(b_n) \right. \\ & \left. + V_\infty(B_n, T - T_n - t)(1 - G(b_n))] dt - V_\infty(B_n, T - T_n) \right\} = 1_{\{T_n \leq T\}} \{ (M_1 V_\infty)(b_n, B_n, T - T_n) - V_\infty(B_n, T - T_n) \}. \end{aligned}$$

Because $V_\infty(B_n, T - T_n) = \sup_{0 \leq b \leq B_n} (M_1 V_\infty)(b, B_n, T - T_n) \geq (M_1 V_\infty)(b_n, B_n, T - T_n)$, (8) implies that $\mathbb{E}[Y_{n+1}^{\mathbf{b}} - Y_n^{\mathbf{b}} | \mathcal{F}_n^{\mathbf{b}}] \leq 0$ for every $n \geq 0$, which proves that $(Y_n^{\mathbf{b}}, \mathcal{F}_n^{\mathbf{b}})_{n \geq 0}$ is a supermartingale. Because

$$\begin{aligned} 1_{\{T_n \leq T\}} (M_1 V_\infty)(b_n^*, B_n, T - T_n) &= 1_{\{T_n \leq T\}} (M_1 V_\infty)(b_\infty(B_n, T - T_n), B_n, T - T_n) \\ &= 1_{\{T_n \leq T\}} \sup_{0 \leq b \leq B_n} (M_1 V_\infty)(b, B_n, T - T_n) = 1_{\{T_n \leq T\}} V_\infty(B_n, T - T_n), \end{aligned}$$

replacing \mathbf{b} with \mathbf{b}^* in (8) gives $\mathbb{E}[Y_{n+1}^{\mathbf{b}^*} - Y_n^{\mathbf{b}^*} \mid \mathcal{F}_n^{\mathbf{b}^*}] = 0$ for every $n \geq 0$, which implies that $(Y_n^{\mathbf{b}^*}, \mathcal{F}_n^{\mathbf{b}^*})_{n \geq 0}$ is a martingale. \square

Below Theorem 3 is the main result and shows that an optimal bidding policy always exists. It establishes the connection of the optimal solution of the dynamic bidding problem to successive approximations in (7) of the unique fixed point of the dynamic programming operator M in (4). This connection later leads to the numerical method explained in the first paragraph of Section 4.

Theorem 3. *For every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$, we have $V(B, T) = V_{\infty}(B, T)$ and that $\mathbf{b}^* = (b_n^*)_{n \geq 0}$ is an optimal bidding strategy for (1), where $b_n^* = b_{\infty}(B_n, T - T_n)$ as in Theorem 1 (v) for every $n \geq 0$.*

Proof. For every $\mathbf{b} = (b_n)_{n \geq 0} \in \mathcal{D}(B, T)$, because $(Y_n^{\mathbf{b}}, \mathcal{F}_n^{\mathbf{b}})_{n \geq 0}$ is a supermartingale by Theorem 2, we have

$$V_{\infty}(B, T) = Y_0^{\mathbf{b}} \geq \mathbb{E}[Y_n^{\mathbf{b}}] = \mathbb{E}\left[\sum_{k=0}^{n-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} + V_{\infty}(B_n, T - T_n) 1_{\{T_n \leq T\}}\right] \geq \mathbb{E}\sum_{k=0}^{n-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}}$$

for every $n \geq 1$, where the last inequality follows from that $V_{\infty}(\cdot, \cdot) \geq 0$. Note that $\lim_{n \rightarrow \infty} \sum_{k=0}^{n-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} = \sum_{k=0}^{\infty} (\mu - b_k) Z_k 1_{\{T_k \leq T\}}$ and $|\sum_{k=0}^{n-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}}| \leq \mu X(T)$ for every $n \geq 1$ and $\mathbb{E}X(T) = \lambda T < \infty$. Therefore, we can take the limit in the last display as $n \rightarrow \infty$, and the dominated convergence theorem implies $V_{\infty}(B, T) \geq \mathbb{E}\sum_{k=0}^{\infty} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} = \mathbb{E}\sum_{k=0}^{\infty} (W_k - b_k) Z_k 1_{\{T_k \leq T\}}$. Taking the supremum of both sides over $\mathbf{b} \in \mathcal{D}(B_0)$ leads to $V_{\infty}(B, T) \geq V(B, T)$ for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$. To show the equality and that $\mathbf{b}^* = (b_n^*)_{n \geq 0} \in \mathcal{D}(B, T)$ is optimal, recall from Theorem 2 that $(Y_n^{\mathbf{b}^*}, \mathcal{F}_n^{\mathbf{b}^*})_{n \geq 0}$ is a martingale, and

$$(9) \quad V_{\infty}(B, T) = Y_0^{\mathbf{b}^*} = \mathbb{E}[Y_n^{\mathbf{b}^*}] = \mathbb{E}\left[\sum_{k=0}^{n-1} (\mu - b_k^*) Z_k 1_{\{T_k \leq T\}} + V_{\infty}(B_n, T - T_n) 1_{\{T_n \leq T\}}\right], \quad n \geq 0.$$

By the dominated convergence, $\lim_{n \rightarrow \infty} \mathbb{E}\sum_{k=0}^{n-1} (\mu - b_k^*) Z_k 1_{\{T_k \leq T\}} = \mathbb{E}\sum_{k=0}^{\infty} (\mu - b_k^*) Z_k 1_{\{T_k \leq T\}}$. Moreover, because $0 \leq V_{\infty}(B_n, T - T_n) 1_{\{T_n \leq T\}} \leq \mu(1 + \lambda(T - T_n)) 1_{\{T_n \leq T\}} \leq \mu(1 + \lambda T) 1_{\{T_n \leq T\}}$ by Proposition 1 (iii),

$$0 \leq \mathbb{E}[V_{\infty}(B_n, T - T_n) 1_{\{T_n \leq T\}}] \leq \mu(1 + \lambda T) \mathbb{P}\{T_n \leq T\} \leq \mu(1 + \lambda T) \mathbb{P}\{X(T) \geq n\} \xrightarrow{n \rightarrow \infty} 0.$$

Therefore, passing to limit in (9) as $n \rightarrow \infty$ leads to $V_{\infty}(B, T) = \mathbb{E}\sum_{k=0}^{\infty} (\mu - b_k^*) Z_k 1_{\{T_k \leq T\}} = \mathbb{E}\sum_{k=0}^{\infty} (W_k - b_k^*) Z_k 1_{\{T_k \leq T\}} \leq V(B, T)$. Because $V_{\infty}(B, T) \geq V(B, T)$ was proved earlier, we conclude that $V_{\infty}(B, T) = V(B, T) = \mathbb{E}\sum_{k=0}^{\infty} (W_k - b_k^*) Z_k$, and that $\mathbf{b}^* \in \mathcal{D}(B, T)$ is an optimal bidding strategy. \square

2.4. The monetary characterization of successive approximations. The next proposition uncovers the economic meaning of $V_n(B, T)$; namely, that $V_n(B, T)$ is the maximum expected net revenue obtained from at most n arrivals of keyword search. As a corollary of the proposition, we shall also establish that $V_n(B, T)$ and $\bar{V}_n(B, T)$ are nondecreasing functions of B, T for every $n \geq 0$. For the formulation and proof of the proposition, let us introduce the processes

$$\begin{aligned} \underline{Y}_{m,0}^{\mathbf{b}} &= \underline{V}_m(B_0, T), \quad \underline{Y}_{m,n}^{\mathbf{b}} := \sum_{k=0}^{n-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} + \underline{V}_{m-n}(B_n, T - T_n) 1_{\{T_n \leq T\}}, \quad 1 \leq n \leq m, \quad m \geq 1, \\ \bar{Y}_{m,0}^{\mathbf{b}} &= \bar{V}_m(B_0, T), \quad \bar{Y}_{m,n}^{\mathbf{b}} := \sum_{k=0}^{n-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} + \bar{V}_{m-n}(B_n, T - T_n) 1_{\{T_n \leq T\}}, \quad 1 \leq n \leq m, \quad m \geq 1 \end{aligned}$$

for every $\mathbf{b} = (b_n)_{n \geq 0} \in \mathcal{D}(B_0)$.

Proposition 2. For every $\mathbf{b} = (b_n)_{n \geq 0} \in \mathcal{D}(B_0)$ and $m \geq 0$, the processes $(\underline{Y}_{m,n}^{\mathbf{b}}, \mathcal{F}_n^{\mathbf{b}})_{n=0}^m$ and $(\bar{Y}_{m,n}^{\mathbf{b}}, \mathcal{F}_n^{\mathbf{b}})_{n=0}^m$ are supermartingales. If $\underline{\mathbf{b}}_m^* = (\underline{b}_{m,n}^*)_{n \geq 0}$ and $\bar{\mathbf{b}}_m^* = (\bar{b}_{m,n}^*)_{n \geq 0}$ are the admissible bidding strategies defined by $\underline{b}_{m,n}^* = \underline{b}_{m-n}(B_n, T - T_n)$ and $\bar{b}_{m,n}^* = \bar{b}_{m-n}(B_n, T - T_n)$ for every $0 \leq n \leq m$, where

$$\underline{b}_n(B, T) = \arg \max_{0 \leq b \leq B} (M_1 \underline{V}_{n-1})(b, B, T) \quad \text{and} \quad \bar{b}_n(B, T) = \arg \max_{0 \leq b \leq B} (M_1 \bar{V}_{n-1})(b, B, T)$$

as in Theorem 1, then for every $m \geq 0$, $(\underline{Y}_{m,n}^{\underline{\mathbf{b}}_m^*}, \mathcal{F}_n^{\underline{\mathbf{b}}_m^*})_{n=0}^m$ and $(\bar{Y}_{m,n}^{\bar{\mathbf{b}}_m^*}, \mathcal{F}_n^{\bar{\mathbf{b}}_m^*})_{n=0}^m$ are martingales. Moreover,

$$\begin{aligned} \underline{V}_m(B, T) &= \sup_{\mathbf{b} \in \mathcal{D}(B, T)} \mathbb{E} \sum_{k=0}^{m-1} (W_k - b_k) Z_k 1_{\{T_k \leq T\}} = \mathbb{E} \sum_{k=0}^{m-1} (W_k - \underline{b}_{m,k}^*) Z_k 1_{\{T_k \leq T\}} \\ (10) \quad \bar{V}_m(B, T) &= \sup_{\mathbf{b} \in \mathcal{D}(B, T)} \mathbb{E} \sum_{k=0}^{m-1} (W_k - b_k) Z_k 1_{\{T_k \leq T\}} + \mathbb{E} [\mu(1 + \lambda(T - T_m)) 1_{\{T_m \leq T\}}] \\ &= \mathbb{E} \sum_{k=0}^{m-1} (W_k - \bar{b}_{m,k}^*) Z_k 1_{\{T_k \leq T\}} + \int_0^T \mu(1 + \lambda(T - t)) \frac{(\lambda t)^{m-1}}{(m-1)!} e^{-\lambda t} \lambda dt \end{aligned}$$

and $\bar{V}_m(B, T) = \underline{V}_m(B, T) + \mathbb{E} [\mu(1 + \lambda(T - T_m)) 1_{\{T_m \leq T\}}]$ for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$ and $m \geq 0$.

Two corollaries follow from Proposition 2. The first corollary, Corollary 1, states that at any time the optimal bidding amount never exceeds the expected potential revenue of a click or the available budget.

Corollary 1. For every $m \geq 0$, each of $\underline{V}_m(B, T)$ and $\bar{V}_m(B, T)$ is nondecreasing in $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$. If the remaining budget equals some $0 \leq B \leq B_{\max}$, then the maximizations in the calculations of successive approximations over bidding amounts can be restricted to from $[0, B]$ to $[0, \mu \wedge B]$; namely,

$$\underline{V}_m(B, T) = \max_{b \in [0, \mu \wedge B]} (M_1 \underline{V}_{m-1})(B, T) \quad \text{and} \quad \bar{V}_m(B, T) = \max_{b \in [0, \mu \wedge B]} (M_1 \bar{V}_{m-1})(B, T).$$

Proof. (10) implies that $\underline{V}_m(B, T)$ and $\bar{V}_m(B, T)$ are nondecreasing in $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$ because i) $\mathcal{D}(B, T)$ grows with B , ii) $X(T)$ is nondecreasing in T , and iii) $T \mapsto \mathbb{E} [\mu(1 + \lambda(T - T_m)) 1_{\{T_m \leq T\}}]$ is increasing. For every $\mu \leq b \leq B$, we have $\mathbb{E}[(W_0 - b)Z_0] = (\mu - b)G(b) \leq 0 = (\mu - \mu)G(\mu) = \mathbb{E}[(W_0 - \mu)Z_0]$ and $\mathbb{E}[\underline{V}_{m-1}(B - bZ_0, T - T_1) 1_{\{T_1 \leq T\}}] \leq \mathbb{E}[\underline{V}_{m-1}(B - \mu Z_0, T - T_1) 1_{\{T_1 \leq T\}}]$ because $\underline{V}_{m-1}(B, T)$ is increasing in $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$. Thus, $\mathbb{E}[(W_0 - b)Z_0 + \underline{V}_{m-1}(B - bZ_0, T - T_1) 1_{\{T_1 \leq T\}}] \leq \mathbb{E}[(W_0 - \mu)Z_0 + \underline{V}_{m-1}(B - \mu Z_0, T - T_1) 1_{\{T_1 \leq T\}}]$ for every $\mu \leq b \leq B$, and

$$\begin{aligned} \underline{V}_m(B, T) &= \max_{0 \leq b \leq B} (M_1 \underline{V}_{m-1})(B, T) = \max_{0 \leq b \leq B} \mathbb{E}[(W_0 - b)Z_0 + \underline{V}_{m-1}(B - bZ_0, T - T_1) 1_{\{T_1 \leq T\}}] \\ &= \max_{0 \leq b \leq \mu \wedge B} \mathbb{E}[(W_0 - b)Z_0 + \underline{V}_{m-1}(B - bZ_0, T - T_1) 1_{\{T_1 \leq T\}}] = \max_{0 \leq b \leq \mu \wedge B} (M_1 \underline{V}_{m-1})(B, T) \end{aligned}$$

for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$, $m \geq 0$. The proof is of the second equation in display is the same. \square

The second corollary of Proposition 2 gives an explicit expression for the difference between the lower and upper bounds on the value function. This upper bound helps us determine the exact number of successive approximations needed for an accurate calculation of the value function.

Corollary 2. Because $\underline{V}_m \uparrow V$ and $\bar{V}_m \downarrow V$ as $m \rightarrow \infty$, for every sufficiently large $m \geq 0$, both \underline{V}_m and \bar{V}_m closely approximate V . In fact, Proposition 2 implies that

$$\underline{V}_m(B, T) \leq V(B, T) \leq \bar{V}_m(B, T) = \underline{V}_m(B, T) + \mathbb{E} [\mu(1 + \lambda(T - T_m)) 1_{\{T_m \leq T\}}]$$

for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$ and $m \geq 0$. Both $\|V - \underline{V}_m\|$ and $\|\bar{V}_m - V\|$ are less than or equal to

$$\begin{aligned} \|\bar{V}_m - \underline{V}_m\| &= \max_{0 \leq T \leq T_{\max}} \mathbb{E}[\mu(1 + \lambda(T - T_m))1_{\{T_m \leq T\}}] = \mathbb{E}[\mu(1 + \lambda(T_{\max} - T_m))1_{\{T_m \leq T_{\max}\}}] \\ &= \int_0^{T_{\max}} \mu(1 + \lambda(T_{\max} - t)) \frac{(\lambda t)^{m-1}}{(m-1)!} e^{-\lambda t} \lambda dt \quad \text{for every } m \geq 0. \end{aligned}$$

Proof of Proposition 2. For every $\mathbf{b} \in \mathcal{D}(B, T)$ and $m \geq 0$, we have $\underline{Y}_{m,n+1}^{\mathbf{b}} - \underline{Y}_{m,n}^{\mathbf{b}}$ equals

$$\begin{aligned} &(\mu - b_n)Z_n 1_{\{T_n \leq T\}} + \underline{V}_{m-n-1}(B_{n+1}, T - T_{n+1})1_{\{T \leq T_{n+1}\}} - \underline{V}_{m-n}(B_n, T - T_n)1_{\{T \leq T_n\}} \\ &= 1_{\{T_n \leq T\}} \left\{ (\mu - b_n)Z_n + \underline{V}_{m-n-1}(B_n - b_n Z_n, T - T_n - (T_{n+1} - T_n))1_{\{T - T_n \leq T_{n+1} - T_n\}} - \underline{V}_{m-n}(B_n, T - T_n) \right\}. \end{aligned}$$

Because the random variables $b_n, B_n, T - T_n$ are $\mathcal{F}_n^{\mathbf{b}}$ -measurable, and $\mathbb{E}[Z_n | \mathcal{F}_n^{\mathbf{b}}] = \mathbb{P}\{Z_n = 1 | b_n\} = G(b_n)$, and $T_{n+1} - T_n$ has exponential distribution with rate λ independently of $\mathcal{F}_n^{\mathbf{b}}$, $\mathbb{E}[\underline{Y}_{m,n+1}^{\mathbf{b}} - \underline{Y}_{m,n}^{\mathbf{b}} | \mathcal{F}_n^{\mathbf{b}}]$ equals

$$\begin{aligned} &\left\{ (\mu - b_n)G(b_n) + \int_0^{T - T_n} \lambda e^{-\lambda t} [\underline{V}_{m-n-1}(B_n - b_n, T - T_n - t)G(b_n) + \underline{V}_{m-n-1}(B_n, T - T_n - t)(1 - G(b_n))] dt \right. \\ &\quad \left. - \underline{V}_{m-n}(B_n, T - T_n) \right\} 1_{\{T_n \leq T\}} = 1_{\{T_n \leq T\}} \left\{ (M_1 \underline{V}_{m-n-1})(b_n, B_n, T - T_n) - \underline{V}_{m-n}(B_n, T - T_n) \right\} \end{aligned}$$

Because $\underline{V}_{m-n}(B, T) = \sup_{0 \leq b \leq B} (M_1 \underline{V}_{m-n-1})(b, B, T) \geq (M_1 \underline{V}_{m-n-1})(b, B, T)$ for every $0 \leq b \leq B$ in general, the display is always nonpositive; namely, $\mathbb{E}[\underline{Y}_{m,n+1}^{\mathbf{b}} - \underline{Y}_{m,n}^{\mathbf{b}} | \mathcal{F}_n^{\mathbf{b}}] \leq 0$ for every $0 \leq n \leq m-1$, and $(\underline{Y}_{m,n}^{\mathbf{b}}, \mathcal{F}_n^{\mathbf{b}})_{n=0}^m$ is a supermartingale for every $\mathbf{b} \in \mathcal{D}(B_0)$. If, however, \mathbf{b} is replaced with $\underline{\mathbf{b}}_m^* = (\underline{b}_{m,n}^*)_{n \geq 0}$ defined by $\underline{b}_{m,n}^* := b_{m-n}(B_n, T - T_n) = \arg \max_{0 \leq b \leq B_n} (M_1 \underline{V}_{m-n-1})(b, B_n, T - T_n)$, then $(M_1 \underline{V}_{m-n-1})(\underline{b}_{m,n}^*, B_n, T - T_n) = \max_{0 \leq b \leq B_n} (M_1 \underline{V}_{m-n-1})(b, B_n, T - T_n) = \underline{V}_{m-n}(B_n, T - T_n)$, and the display equals identically zero. Thus, $\mathbb{E}[\underline{Y}_{m,n+1}^{\underline{\mathbf{b}}_m^*} - \underline{Y}_{m,n}^{\underline{\mathbf{b}}_m^*} | \mathcal{F}_n^{\underline{\mathbf{b}}_m^*}] = 0$ for every $0 \leq n \leq m-1$, and $(\underline{Y}_{m,n}^{\underline{\mathbf{b}}_m^*}, \mathcal{F}_n^{\underline{\mathbf{b}}_m^*})_{n=0}^m$ is a martingale. The supermartingale property implies

$$\begin{aligned} \underline{V}_m(B_0, T) = \underline{Y}_{m,0}^{\mathbf{b}} &\geq \mathbb{E}[\underline{Y}_{m,m}^{\mathbf{b}}] = \mathbb{E} \left[\sum_{k=0}^{m-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} + \underline{V}_0(B_m, T - T_m) 1_{\{T_m \leq T\}} \right] \\ &= \mathbb{E} \left[\sum_{k=0}^{m-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} \right] \quad \text{for every } \mathbf{b} = (b_n)_{n \geq 0} \in \mathcal{D}(B_0), \end{aligned}$$

because $\underline{V}_0 \equiv 0$. Taking the sup of both sides over all $\mathbf{b} \in \mathcal{D}(B_0)$ gives $\underline{V}_m(B_0, T) \geq \sup_{\mathbf{b} \in \mathcal{D}(B_0)} \mathbb{E}[\sum_{k=0}^{m-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}}]$. However, if we replace \mathbf{b} with $\underline{\mathbf{b}}_m^*$, then the inequality in the last display becomes an equality because $(\underline{Y}_{m,n}^{\underline{\mathbf{b}}_m^*}, \mathcal{F}_n^{\underline{\mathbf{b}}_m^*})_{n=0}^m$ is a martingale, and we obtain $\underline{V}_m(B_0, T) = \mathbb{E}[\sum_{k=0}^{m-1} (\mu - \underline{b}_{m,k}^*) Z_k 1_{\{T_k \leq T\}}] \leq \sup_{\mathbf{b} \in \mathcal{D}(B_0)} \mathbb{E}[\sum_{k=0}^{m-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}}]$. The last two opposite inequalities imply the equalities on the first line of (10) because the common mean μ can be replaced with unobservable i.i.d. random variables W_0, W_1, \dots , which are independent of the filtration $(\mathcal{F}_n^{\mathbf{b}})_{n \geq 0}$ for every $\mathbf{b} \in \mathcal{D}(B_0)$.

One can similarly prove that $(\bar{Y}_{m,n}^{\mathbf{b}}, \mathcal{F}_n^{\mathbf{b}})_{n=0}^m$ is a supermartingale for every $\mathbf{b} \in \mathcal{D}(B_0)$ and $(\bar{Y}_{m,n}^{\bar{\mathbf{b}}_m^*}, \mathcal{F}_n^{\bar{\mathbf{b}}_m^*})_{n=0}^m$ is a martingale. The supermartingale property implies

$$\begin{aligned} \bar{V}_m(B_0, T) = \bar{Y}_{m,0}^{\mathbf{b}} &\geq \mathbb{E}[\bar{Y}_{m,m}^{\mathbf{b}}] = \mathbb{E} \left[\sum_{k=0}^{m-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} + \bar{V}_0(B_m, T - T_m) 1_{\{T_m \leq T\}} \right] \\ &= \mathbb{E} \left[\sum_{k=0}^{m-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} \right] + \mathbb{E}[\mu(1 + \lambda(T - T_m))1_{\{T_m \leq T\}}] \end{aligned}$$

for every $\mathbf{b} = (b_n)_{n \geq 0} \in \mathcal{D}(B_0)$, because $\bar{V}_0(B, T) = \mu(1 + \lambda T)$ for every $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$ as in (7). If \mathbf{b} is replaced with $\bar{\mathbf{b}}_m^*$, then the inequality in the display holds with equality. Therefore, $\bar{V}_m(B_0, T) = \sup_{\mathbf{b} \in \mathcal{D}(B_0)} \mathbb{E}[\sum_{k=0}^{m-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}}] + \mathbb{E}[\mu(1 + \lambda(T - T_m))1_{\{T_m \leq T\}}] = \mathbb{E}[\sum_{k=0}^{m-1} (\mu - \bar{b}_{m,k}^*) Z_k 1_{\{T_k \leq T\}}] +$

$\int_0^T \mu(1 + \lambda(T-t))(\lambda t)^{m-1} \lambda dt / (m-1)!$, because the m th-arrival time T_m of Poisson process X has Erlang distribution with parameters λ and m . This proves the remainder of (10) because W_0, W_1, \dots are i.i.d. with mean μ independent of $X, B_0, B_1, \dots, Z_0, Z_1, \dots$. \square

2.5. The maximum expected total net revenue under optimal bidding policy. Recall that, as we defined the value function $V(\cdot, \cdot)$ we assumed that the first keyword search arrives at time 0 and the others according to Poisson process X afterwards; this assumption was necessary for evoking the dynamic programming optimality principle in (2). In reality, however, the arrival time of the first keyword search is strictly positive with probability one and has exponential distribution with mean $1/\lambda$ because we assumed that keyword search instances arrive according to a Poisson process with arrival rate λ . Let us denote the maximum expected total net revenue obtainable with budget B over time horizon T by $U(B, T)$. Then

$$(11) \quad U(B, T) = \mathbb{E}[V(B, T - T_1)1_{\{T_1 \leq T\}}] = \int_0^T \lambda e^{-\lambda t} V(B, T - t) dt, \quad (B, T) \in [0, B_{\max}] \times [0, T_{\max}].$$

Because $V(B, T) \geq 0$ is jointly increasing in B and T , we have $U(B, T) \leq \mathbb{E}[V(B, T)1_{\{T_1 \leq T\}}] \leq V(B, T)$; however, the difference $V(B, T) - U(B, T) > 0$ decreases with increasing keyword arrival rate λ .

2.6. Finding the value function of any admissible Markovian bidding policy. Beside the optimal bidding strategy found by solving the associated dynamic programming problem, we shall numerically study the performances of some other reasonable bidding strategies, introduced in Section 3. Those include the dynamical versions of the optimal bidding strategy “BC” of Cholette, et al. [1] under the “expected total cost not to exceed the budget” constraint and the optimal strategy which greedily bids the same amount until either time expires or the budget is depleted, whichever occurs first. Because the latter two strategies have significantly less running-time and space requirements than those of the dynamic-programming approach, they can be considered as some viable alternatives as the numerical examples in the next section suggest. Here we shall explain how one can calculate the value function of an admissible bidding strategy whose bidding amount at any time depends only on the remaining budget and remaining time at that moment.

Suppose that $V^p(B, T)$ denotes the value function of an admissible policy p which bids some known amount $0 \leq b(B, T) \leq \mu \wedge B$ when budget B is left for the remaining T units of time; namely,

$$V^p(B, T) = \mathbb{E} \sum_{i=0}^{\infty} [W_i - b(B_i, T - T_i)] Z_i 1_{\{T_i \leq T\}} \quad \text{for every } (B, T) \in [0, B_{\max}] \times [0, T_{\max}],$$

where $T_0 = 0, T_1, T_2, \dots$ are the arrival times of Poisson keyword search process, $B_i = B_{i-1} - b(B_{i-1}, T_{i-1}) Z_{i-1}$ is the budget left after $(i-1)$ st keyword search, Z_i is the indicator of a click which happens with probability $G(b(B_i, T_i))$ for every $i \geq 0$. Below we state some useful facts about $V^p(B, T)$. Their proofs are omitted because they are very similar to those for the facts stated earlier for the value function $V(B, T)$. First of all,

$$V^p(B, T) = \mathbb{E}[(W_0 - b(B, T))Z_0 + V^p(B - b(B, T)Z_0, T - T_1)1_{\{T_1 \leq T\}}] \equiv (M^p V^p)(B, T),$$

where

$$\begin{aligned} (M^p w)(B, T) &:= (M_1 w)(b(B, T), B, T) = (\mu - b(B, T))G(b(B, T)) \\ &\quad + \int_0^T \lambda e^{-\lambda t} [w(B - b(B, T), T - t)G(b(B, T)) + w(B, T - t)(1 - G(b(B, T)))] dt \end{aligned}$$

is a contraction operator on the collection of bounded functions $\mathcal{B}([0, B_{\max}] \times [0, T_{\max}])$ on $[0, B_{\max}] \times [0, T_{\max}]$ with $\|M^p w_1 - M^p w_2\| \leq (1 - e^{-\lambda T_{\max}})\|w_1 - w_2\|$. We have $0 \leq V^p(B, T) \leq \mu(1 + \lambda T)$ for every $(B, T) \in$

$[0, B_{\max}] \times [0, T_{\max}]$. The value function V^p is the unique bounded fixed point of M^p ; therefore, $V^p = M^p V^p$. The lower and upper successive approximations

$$\begin{aligned} \underline{V}_0^p(B, T) &\equiv 0 & \text{and} & \quad \underline{V}_n^p(B, T) := (M^p \underline{V}_{n-1}^p)(B, T), \quad (B, T) \in [0, B_{\max}] \times [0, T_{\max}], \\ \overline{V}_0^p(B, T) &= \mu(1 + \lambda T) & \text{and} & \quad \overline{V}_n^p(B, T) := (M^p \overline{V}_{n-1}^p)(B, T), \quad (B, T) \in [0, B_{\max}] \times [0, T_{\max}] \end{aligned}$$

are increasing and decreasing, respectively, and converge pointwise to $V^p(B, T)$ as $n \rightarrow \infty$, uniformly in $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$ and

$$\begin{aligned} \|V^p - \underline{V}_n^p\| &\leq \mu(1 + \lambda T_{\max}) e^{\lambda T_{\max}} (1 - e^{-\lambda T_{\max}})^n \quad \text{for every } n \geq 0, \\ \|V^p - \overline{V}_n^p\| &\leq \mu(1 + \lambda T_{\max}) e^{\lambda T_{\max}} (1 - e^{-\lambda T_{\max}})^n \quad \text{for every } n \geq 0. \end{aligned}$$

The process

$$Y_0^p := V^p(B_0, T) \quad \text{and} \quad Y_n^p := \sum_{k=0}^{n-1} (\mu - b_k) Z_k 1_{\{T_k \leq T\}} + V^p(B_n, T - T_n) 1_{\{T_n \leq T\}}, \quad n \geq 0$$

is a *martingale* adapted to the filtration $(\mathcal{F}_n^p)_{n \geq 0}$ defined by $\mathcal{F}_n^p = \sigma\{B_0, T_0, \dots, B_n, T_n\}$ for every $n \geq 0$. If we denote the expected total net revenue under policy p by U^p , then we have

$$U^p(B, T) = \mathbb{E}[V^p(B, T - T_1) 1_{\{T_1 \leq T\}}] = \int_0^T \lambda e^{-\lambda t} V^p(B, T - t) dt, \quad (B, T) \in [0, B_{\max}] \times [0, T_{\max}].$$

3. ALTERNATIVE STATIC AND DYNAMIC BIDDING STRATEGIES

Optimal bidding strategy can be found by solving the dynamic programming problem in (1), which may require excessive computational time and space for large values of parameters $\mu, \lambda, B_{\max}, T_{\max}$. Here, we will propose one static and two dynamic bidding strategies, whose computations are less demanding, and compare their performances with that of the optimal bidding strategy on a large number of numerical examples in Section 4. Our solution of the optimal bidding problem in Section 2 thus not only provides an optimal bidding policy but also enables us to search for simple and good-performing alternative bidding policies for the cases when optimal bidding policies are hard to calculate.

3.1. Static and dynamic budget-constrained bidding strategy. Cholette et al. [1] formulate the so-called *budget-constrained (BC) bidding* problem

$$\begin{aligned} \overline{U}^{BC}(B, T) &= \max_b \mathbb{E} \left[\sum_{i=1}^{\infty} (W_i - b) Z_i 1_{\{T_i \leq T\}} \right] \text{ subject to } \mathbb{E} \left[\sum_{i=1}^{\infty} b Z_i 1_{\{T_i \leq T\}} \right] \leq B \\ &= \max_b (\mu - b) \lambda T G(b) \quad \text{subject to} \quad b \lambda T G(b) \leq B. \end{aligned}$$

Let us denote by $b^{BC}(B, T)$ the bid amount b that attains the maximum and call it *BC-optimal* bid amount when budget B is left for the remaining T units of time. The budget-constrained bidding problem can be solved with the Lagrange-relaxation method in general. The problem admits closed-form solution

$$b^{BC}(B, T) = \begin{cases} \sqrt{a^2 + a\mu} - a, & \frac{\lambda T (\sqrt{a^2 + a\mu} - a)^2}{\sqrt{a^2 + a\mu}} \leq B \\ \frac{B + \sqrt{B^2 + 4\lambda T B a}}{\lambda T}, & \text{otherwise} \end{cases}, \quad \text{if } m = 1, p_0 = 1, \text{ and } p_1 = 0.$$

The original formulation of Cholette, et al. [1] assumes that the same bidding amount $b = b^{BC}(B, T)$ is used until either the budget is depleted or time is up, whichever occurs first. However, one expects to do better than $U^{BC}(B, T)$ by recalculating the BC-optimal bid amount at each keyword search instance in the future.

Here we propose the *dynamic budget-constrained (DBC)* bidding strategy in the following way: at the arrival times T_1, T_2, T_3, \dots of keyword search instances, let us bid in the amount $b^{BC}(B_0 \equiv B, T), b^{BC}(B_1, T - T_1), b^{BC}(B_2, T - T_2), \dots$, respectively, as long as we have both budget and time left. We shall denote the expected total net revenue under DBC policy by U^{DBC} , which can be calculated with successive approximations described in Section 2.6.

For a fair comparison of (static) BC policy with other dynamic bidding strategies, we shall also calculate its expected total net revenue

$$U^{BC}(B, T) = (\mu - b^{BC}(B, T)) \sum_{i=1}^{\lfloor B/b^{BC}(B, T) \rfloor} \left(1 - \sum_{j=0}^{i-1} \frac{(\lambda G(b^{BC}(B, T))T)^j}{j!} e^{-\lambda G(b^{BC}(B, T))T} \right)$$

under *strict* budget-constraint. The derivation follows from thinning the Poisson arrivals of keyword search instances, as discussed in detail for the calculation of the expected total net revenue of the static greedy bidding policy of the next section.

3.2. Static and dynamic greedy bidding strategy. If one wants to bid the same amount b all the time, then the budget constraint can be forced not just in expectation but with probability one.

Note in this case that $(Z_i, T_i)_{i=1}^\infty$ form a marked Poisson process, where the marks Z_1, Z_2, \dots are i.i.d. Bernoulli random variables with common success probability $G(b)$. If the $(S_i)_{i=1}^\infty$ are those arrival times T_i corresponding to $Z_i = 1$, then the $(S_i)_{i=0}^\infty$ turn out to be the arrival times of a Poisson process with arrival rate $\lambda G(b)$. At each arrival time S_i , the budget is decreased by the amount b , in return of $W_i - b$ net revenue. Thus, starting with budget B for the remaining T units of time, one can generate $\lfloor B/b \rfloor$ or as many as the number of arrivals $(S_i)_{i=1}^\infty$ in $(0, T]$, whichever is the smaller. Let us call an amount b that attains the maximum in

$$U^G(B, T) = \max_b \mathbb{E} \left[\sum_{i=1}^{\lfloor B/b \rfloor} (W_i - b) 1_{\{S_i \leq T\}} \right]$$

the (static) *greedy bidding* strategy and denote it by $b^G(B, T)$, where $\lfloor x \rfloor$ is the largest integer smaller than or equal to x . Because $Y(t) := \sum_{i=1}^\infty 1_{\{S_i \leq t\}}$, $t \geq 0$ is a Poisson process with rate $\lambda G(b)$, and

$$\mathbb{E} \left[\sum_{i=1}^{\lfloor B/b \rfloor} (W_i - b) 1_{\{S_i \leq T\}} \right] = (\mu - b) \sum_{i=1}^{\lfloor B/b \rfloor} \mathbb{P}\{S_i \leq T\} = (\mu - b) \sum_{i=1}^{\lfloor B/b \rfloor} \mathbb{P}\{Y(T) \geq i\},$$

the greedy bidding problem reduces to

$$U^G(B, T) = \max_b (\mu - b) \sum_{i=1}^{\lfloor B/b \rfloor} \left(1 - \sum_{j=0}^{i-1} \frac{(\lambda G(b)T)^j}{j!} e^{-\lambda G(b)T} \right).$$

We again expect that recalculating the greedy bidding amount before every keyword search will improve the performance. Therefore, we also propose the following *dynamic greedy (DG)* bidding strategy: at the arrival times T_1, T_2, T_3, \dots of keyword search instances, let us bid in the amount $b^G(B_0 \equiv B, T), b^G(B_1, T - T_1), b^G(B_2, T - T_2), \dots$, respectively, as long as we have both budget and time left. We shall denote the expected total net revenue under DG policy by U^{DG} , which can be calculated with successive approximations described in Section 2.6.

In the meantime, note that we always have $U^G(B, T) \geq U^{BC}(B, T)$. Hence, static greedy bidding policy should perform at least as good as the static budget-constraint bidding policy if the budget-constraint is strictly enforced. However, the numerical examples in Section 4 suggest that the difference $U^G(B, T) - U^{BC}(B, T) \geq 0$ is often surprisingly small.




FIGURE 4. The convergence rate of the dynamic programming algorithm

4. NUMERICAL EXAMPLES

As Corollary 2 suggests, one can calculate iteratively the sequences of decreasing and increasing successive approximations, respectively $(\bar{V}_n)_{n \geq 0}$ and $(\underline{V}_n)_{n \geq 0}$ in (7), until the first $n_\varepsilon \geq 1$ such that $\|\bar{V}_{n_\varepsilon} - \underline{V}_{n_\varepsilon}\| \leq \varepsilon$ for some sufficiently small $\varepsilon > 0$. We can then take, for example, $V \approx \bar{V}_{n_\varepsilon}$, because $\underline{V}_n \leq V \leq \bar{V}_n$ for every $n \geq 1$ by Theorem 1 (ii & iv) and Theorem 3, and $\|\bar{V}_{n_\varepsilon} - V\| \leq \|\bar{V}_{n_\varepsilon} - \underline{V}_{n_\varepsilon}\| \leq \varepsilon$. Because no closed form expressions of \bar{V}_n or \underline{V}_n for $n \geq 1$ exist, we calculate them on a fine grid of their common domain $[0, B_{\max}] \times [0, T_{\max}]$ of the Cartesian product of remaining budget and remaining time. Since the remaining budget is discretized, the bidding amounts are also discretized and equal to some integer multiple of grid size used for the discretization of remaining budget. Let us also remember from Corollary 1 that optimal bidding amounts are always less than or equal to expected sales revenue $\mu = \mathbb{E}W_0$.

We use the same parameters; namely, $a = 20$, $\lambda = 500$ keyword-search per day, $T_{\max} = 1$ day, $m = 1$, $p_0 = 1$, $p_1 = 0$, $B_{\max} = \$3000$ per day, $\mu = \$50$ per click, as in the single keyword example of Cholette et. al. [1, Section 5.5]. The intervals $[0, B_{\max}]$ and $[0, T_{\max}]$ are divided into 3000 and 300 subintervals of equal length \$1 and 1/300 day, respectively. The optimal bidding amounts are searched over those values in $\{\$0, \$1, \$2, \dots, \$49, \$50\}$ which are smaller than or equal to the remaining budget.

The explicit theoretical upperbound $\|\bar{V}_m - \underline{V}_m\|$, given by Corollary 2, on the difference between successive approximations $\bar{V}_m, \underline{V}_m$ and value function V is evaluated and plotted in Figure 4. The approximation error sharply decreases after 450 iterations and becomes less than a penny by the 600th iteration.

We do not depend on the theoretical error bound of Corollary 2 and run the successive approximations until the maximum difference between upper and lower bounds on value function reduces below one cent. This takes about 620 iterations, which is consistent with the reading from Figure 4. The value function $V(B, T)$ for $(B, T) \in [0, B_{\max}] \times [0, T_{\max}]$ and its contour lines are plotted on top in Figure 5. Even though the value function equals $V(T_{\max}, B_{\max}) = V(1, 3000) = \7420.98 , the maximum expected total net revenue is slightly less than this. Remember from (11) that the maximum expected total net revenue is $U(B, T) = \mathbb{E}[V(B, T - T_1)1_{\{T_1 \leq T\}}] \leq V(B, T)$. However, because $\lambda = 500$ is large, the first keyword search arrives rather soon with high probability, and U and V are not expected to be very different. Indeed, the

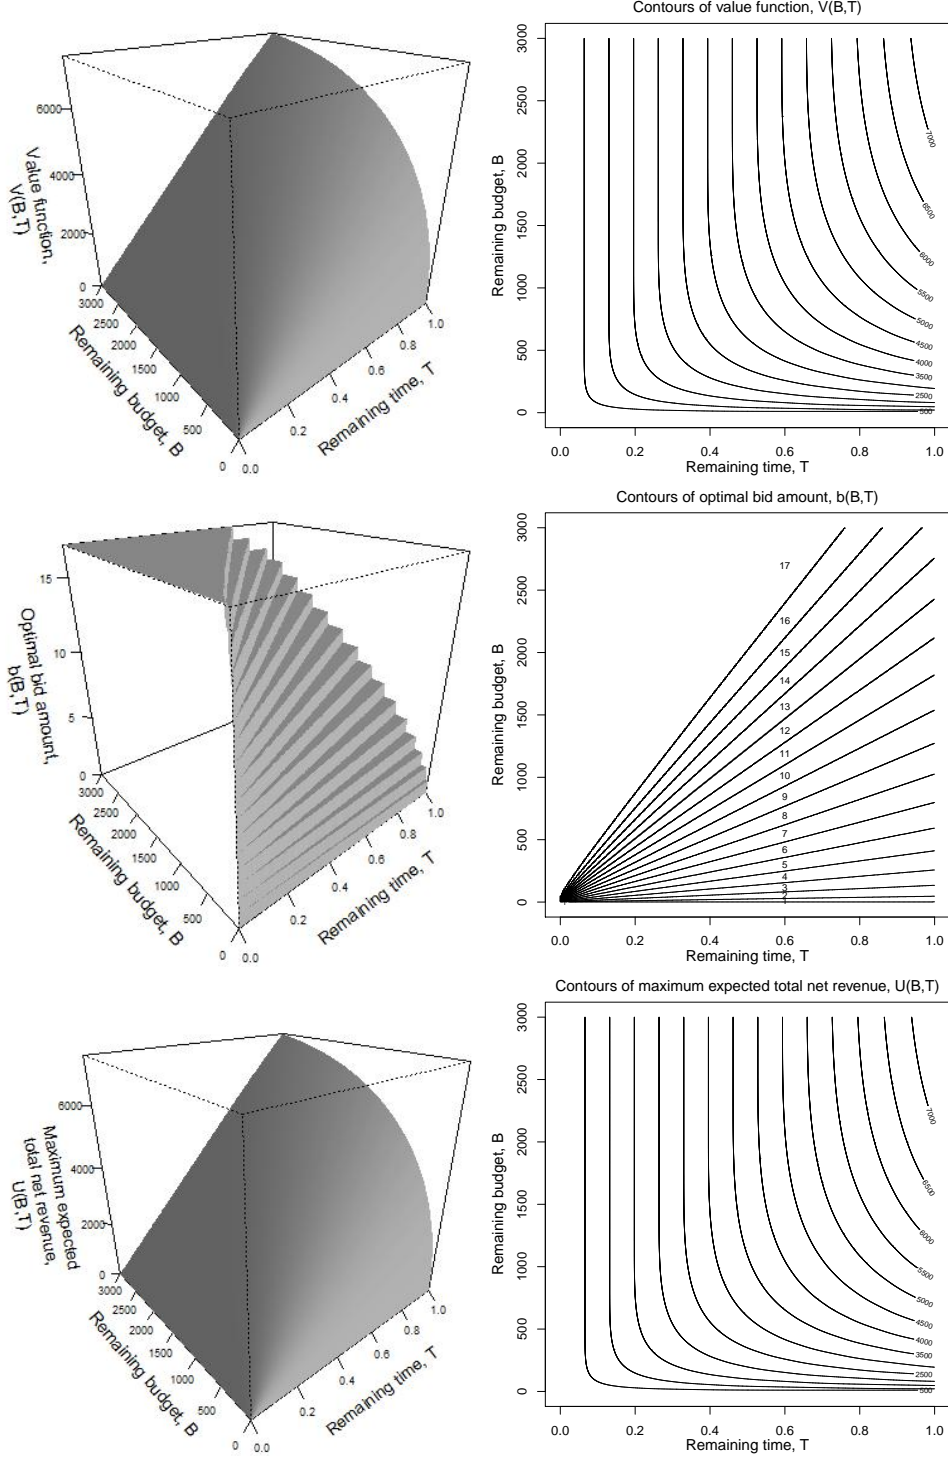


FIGURE 5. First two rows: the value function $V(B, T)$ and optimal policy function $b(B, T)$ of the dynamic programming problem. As the grid size used for the discretization of the remaining budget decreases, the ladder-like optimal bidding policy function converges to a smooth surface cut from top by a horizontal plane. Last row: the maximum expected total net revenue $U(B, T)$ calculated by (11).

Methods	No constraints (NC)	Budget constraint (BC)	Probabilistic constraint (PC)	Trade-off solution (TO)	Dynamic strategy (DS)
Bid amount	17.4	14.3	4.97	7.0	dynamic ^a
Expected total net revenue	7583.4	7447.3	4482.1	5574.1	7407.85
Expected total cost	4053.5	3000	494.6	907.4	2986 ^b
Overbudget probability	0.53	0.49	0.10	0.25	0
Expected total net revenue under strict budget constraint ^c	5607.20	7245.79	4481.36	5574.07	7407.85
Expected total cost under strict budget constraint ^d	2992.80	2902.37	494.61	907.41	2986 ^b

^a Optimal bidding policy is plotted in the second row of Figure 5.

^b Estimated by simulation: 1000 replications give mean 2986 with 0 (zero) standard error.

^{c,d} Respectively, $(\mu - b)\mathbb{E}(Y \wedge \lfloor \frac{B}{b} \rfloor)$ and $b\mathbb{E}(Y \wedge \lfloor \frac{B}{b} \rfloor)$, where $Y \sim \text{Pois}(\lambda G(b)T)$; $T = 1$ and $B = 3000$ are the initial remaining time and budget, respectively; b is the bid amount in the first row.

TABLE 1. Comparison of optimal static and dynamic bidding strategies. NC, BC, PC, TO are optimal static strategies proposed by Cholette et al. [1]. Their performance measures (except ^{c,d}) are copied from Table 2 in [1]. The last column shows the performance of the dynamic bidding strategy proposed here.

plots of V and U in Figure 5 look indifferent, but it turns out $U(1, 3000) = 7407.85 < 7420.98 = V(1, 3000)$.

4.1. Comparison of optimal static and dynamic bidding strategies. In Table 1, we compare the performance measures of the optimal dynamic bidding strategy (DS) with those of the optimal static strategies (NC, BC, PC, TO) proposed by Cholette et al. [1]. For all of those static bidding strategies the budget constraint is soft. NC assumes no budget constraints at all, BC requires that the expected total cost be exactly 3000, PC wants the probability that the total costs exceeds 3000 is less than or equal to 0.10, and TO picks an arbitrary point on the expected profit/probability efficient frontier and calculates the corresponding bid amount. According to the “maximum total net revenue” measure in the second line of Table 1, DS performs slightly under NC and BC, while NC incurs on average $(4053.5 - 3000)/3000 = 35\%$ more cost than available budget and both BC and NC are likely to go over budget nearly 50% of the time. Even though PC and TO violate the budget constraint much less frequently (10% and 25%, respectively) and use only a small fraction (on average, $494.6/3000 = 14\%$ and $907.4/3000 = 30\%$) of the available budget, they are outperformed by DS by a large margin. For a more fair comparison, the last two rows of Table 1 report the same quantities (the expected total net revenue and total cost) under the *strict* budget constraint. The performances of overly cautious PC and TO do not change, but both the expected total net revenues and total costs of NC and BC are marked down. Under the strict budget constraint, DS has an expected total net revenue $(7407.85 - 5607.2)/5607.2 = 32\%$ and $(7407.85 - 7245.79)/7407.85 = 2\%$ more than those of NC and BC, respectively, while incurring an expected total cost $(2992.80 - 2986)/2992.80 = 0.2\%$ less and $(2986 - 2902.37)/2902.37 = 2.9\%$ more than those of NC and BC, respectively. When the budget constraint is strictly enforced, the expected total net revenue is the sole performance measure, according to which DS outperforms all of the static strategies as expected.

The optimal bidding policy is a function of only the remaining time and remaining budget. It is displayed in the second row of Figure 5. The ladder-like piecewise constant look is an artifact due to the discretization of remaining-budget variable and disappears as the grid size decreases to zero. In the limit, the optimal

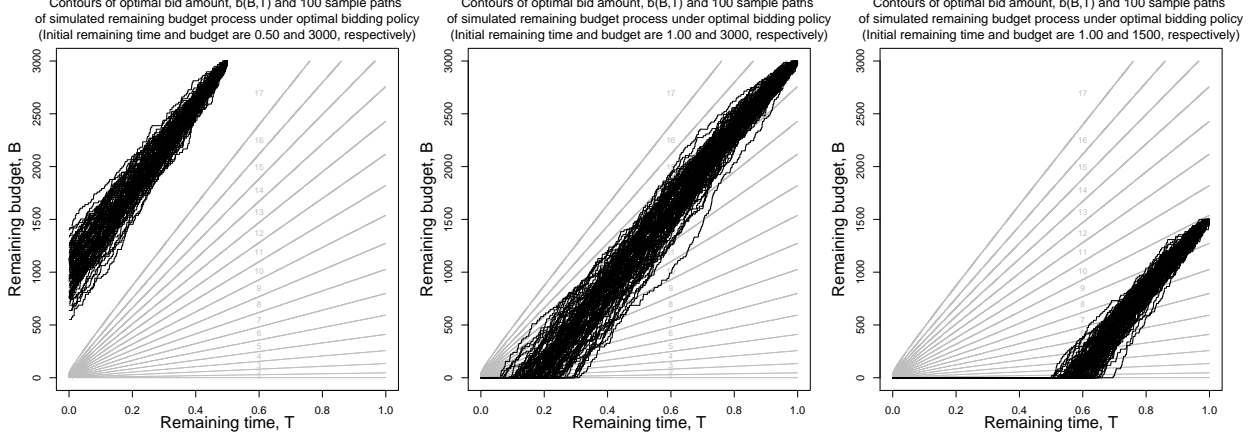


FIGURE 6. Simulation of remaining budget processes under optimal bidding policy for different initial remaining time and remaining budget values

bidding policy is a smooth surface cut from top by a flat plane. For a fixed T , optimal bid amount increases as B increases. Also, for a fixed B , optimal bid amount decreases as T increases.

Figure 6 shows, from left to right, 100 realizations of the remaining-budget process under optimal bidding policy if the initial remaining time and budget are $(0.5, 3000)$, $(1.0, 3000)$, and $(1.0, 1500)$, respectively. If no keyword-search instances arrive for relatively long time periods, then the remaining budget process remains flat over time and enters into the higher-bidding-regions and optimal bidding amounts increase. If the keyword-search instances arrive relatively quickly, then the remaining budget process may drift downward into lower-bidding-regions and optimal bidding amount decrease. It is interesting that the slope of the “average behavior” of remaining-budget process is the same for all starting points at time zero. The optimal bidding strategy has to balance the increasing (respectively, decreasing) likelihood of inducing an actual click with the increasing (respectively, decreasing) cost of frequent high bids (respectively, low bids). This trade-off seems to be optimally balanced by keeping the direction of the process around an ideal constant, which is independent of the remaining time and budget at any time.

4.2. Sensitivity analysis. The performances of static budget-constraint and greedy bidding policies (BC and G, respectively) and the dynamic budget-constraint and greedy bidding policies (DBC and DG, respectively) are compared with that of optimal dynamic bidding policy (DS) for five different levels of each parameter value B, a, λ, m, μ while the others are kept fixed. Table 2 summarizes the results. Recall from Section 3.1 that \bar{U}^{BC} is the maximum expected total net revenue under the budget constraint on the *expected* total cost. Therefore, we did not compare it directly to other static and dynamic bidding strategies, which are defined under strict (with probability one) budget constraint.

The relative losses with respect to optimal dynamic bidding policy are the largest (respectively, 2.35-3.26% and 1.28-1.74%) when the budget is the tightest (750-1500) and the arrival rate is the highest (1000-2000). The dynamic budget-constraint (DBC) bidding policy is superior to the static budget-constraint (BC) bidding policy in general by an average margin of 1.5%. Similarly, the dynamic greedy (DG) bidding policy performs better than the static greedy (G) bidding policy in general by an average margin of 0.3%. Hence, the static greedy bidding policy may serve as a good replacement for its dynamic counterpart. The dynamic greedy bidding policy outperforms the dynamic budget-constraint bidding policy by a small margin in general. In almost all cases, the relative losses of the dynamic greedy and budget-constraint bidding policies are less than 1% and provide near-optimal expected total net revenue at significantly lower computational requirements than those of the dynamic programming approach.

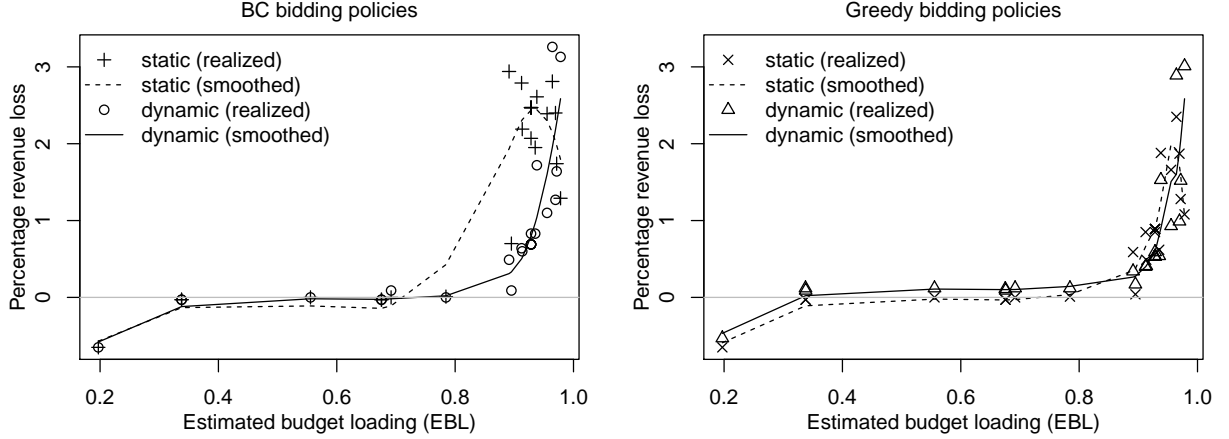


FIGURE 7. The comparison of the static and dynamic bidding policies

Table 2 suggests that the larger is the initial budget relative to the expected total spending on bidding, the better are the performances of the static/dynamic BC/greedy bidding policies. To quantify this relation and understand it better, we introduce a surrogate for the fraction of the initial budget spent on bidding; namely,

$$\text{estimated budget loading (EBL)} = \frac{b^G(B, T) \lambda G(b^G(B, T)) T}{B},$$

where $b^G(B, T)$ is the optimal (static) greedy bidding amount of Section 3.2 and the numerator is an estimate of the optimal expected total spending for bidding under optimal dynamic bidding policy. The latter quantity should ideally be calculated by solving the dynamic programming problem. This would however defeat our purpose of finding a fast and accurate alternative bidding policy in place of the optimal dynamic bidding policy when the latter is difficult to calculate.

For each row of Table 2, we calculate EBL (using b^G column) and plot the relative percentage revenue loss of each bidding policy (static/dynamic BC/greedy) in total expected net revenues (columns $U^{BC}, U^G, U^{DBC}, U^{DG}$) against EBL in Figure 7. To better visualize the relation between the relative loss and EBL, we also add to the plots the curves obtained by applying loess (locally weighted scatterplot smoothing) to the data points. Four smoothed curves corresponding to static/dynamic BC/greedy bidding policies were also plotted together in Figure 2.

The negative percentage revenue loss values at $\text{EBL} = 0.2$ in pictures of Figure 7 and 2 are the artifacts of that the bidding amounts are restricted to the one-dollar increments as we numerically solve the stochastic dynamic optimization problem. When the bidding amounts are searched over the one-cent increments, those artifacts disappear; however, the grid size increases 100 folds and Table 2 turns out to be impossible to calculate in a reasonable amount of time. This is known as the dynamic programming's curse of dimensionality, which is exactly the reason why we proposed and studied alternative static/dynamic BC/greedy bidding policies.

All four (static/dynamic BC/greedy) bidding policies perform nearly as good as the optimal dynamic bidding policy for low to moderately-high budget loading values: the percentage revenue losses are nearly zero for BC bidding policies when $\text{EBL} \leq 0.7$ and for static greedy and dynamic BC/greedy policies when $\text{EBL} \leq 0.8 \sim 0.9$. As the estimated budget loading (EBL) increases to one, the percentage revenue losses sharply increase. Interestingly, the relative percentage losses for the static BC/greedy policies seem to be

Parameters					Static bidding values		\bar{U}^{BC}	Relative loss (percentage) with respect to U^{DS} in				$U^{DS\dagger}$
B	a	λ	m	μ	b^{BC}	b^G		U^{BC}	U^G	U^{DBC}	U^{DG}	
750					6.27	6.15	5222.92	2.81	2.35	3.26	2.89	5167.64
1500					9.38	9.04	6487.30	2.61	1.88	1.72	1.53	6435.35
3000	20	500	1	50	14.35	13.70	7447.27	2.47	0.89	0.69	0.53	7407.85
6000					17.41	17.42	7583.42	-0.03	-0.03	-0.03	0.09	7581.09
12000					17.41	17.42	7583.42	-0.03	-0.03	-0.03	0.09	7581.09
	5				9.24	8.80	13224.98	1.95	0.62	0.83	0.54	13178.40
	10				11.30	10.75	10266.55	2.07	0.85	0.83	0.59	10220.10
3000	20	500	1	50	14.35	13.70	7447.27	2.47	0.89	0.69	0.53	7407.85
	40				18.77	17.54	4987.33	2.94	0.59	0.49	0.34	4964.06
	80				21.98	21.90	3019.60	0.01	0.01	0.00	0.12	3019.52
		125			17.41	17.41	1895.85	-0.03	-0.03	-0.03	0.12	1895.27
		250			17.41	17.41	3791.71	-0.03	-0.03	-0.03	0.11	3790.55
3000	20	500	1	50	14.35	13.70	7447.27	2.47	0.89	0.69	0.53	7407.85
		1000			9.38	9.23	12974.66	1.74	1.28	1.64	1.52	12899.30
		2000			6.27	6.20	20891.68	1.29	1.08	3.13	3.01	20759.90
			0.25		7.97	7.50	15223.74	0.70	0.04	0.09	0.17	15218.10
			0.5		10.35	9.68	11485.91	2.19	0.49	0.60	0.41	11452.80
3000	20	500	1	50	14.35	13.70	7447.27	2.47	0.89	0.69	0.53	7407.85
			2		21.69	20.69	3914.34	2.79	0.85	0.64	0.40	3886.07
			4		29.88	29.88	1397.46	-0.00	-0.00	0.09	0.12	1397.42
				12.5	5.49	5.49	754.90	-0.65	-0.65	-0.65	-0.53	750.00
				25	10.00	10.00	2500.00	0.00	0.00	0.00	0.12	2500.00
3000	20	500	1	50	14.35	13.70	7447.27	2.47	0.89	0.69	0.53	7407.85
				100	14.35	13.95	17894.54	2.39	1.66	1.10	0.93	17785.20
				200	14.35	14.08	38789.08	2.40	1.87	1.27	0.99	38554.60

[†] Because the maximum bid amount is searched over integers, the maximum expected total net revenue calculations involve a small negative discretization bias, which explains those negative relative losses in the table.

TABLE 2. Sensitivity analysis. The relative loss of U^{BC} with respect to U^{DS} is defined by $100(U^{DS} - U^{BC})/U^{DS}$, and the others are defined likewise.

bounded, while the relative percentage losses of the dynamic BC/greedy policies seem to grow unboundedly as the budget loading increases to one. Out of four bidding policies proposed as an alternative to the optimal dynamic bidding policy, the static greedy bidding policy seems to be the best.

ACKNOWLEDGMENTS

The numerical calculations reported in this paper were performed at TUBITAK ULAKBIM, High Performance and Grid Computing Center (TR-Grid e-Infrastructure). The research of Savas Dayanik was partially supported with the TUBITAK Research Grant 110M610. Mahmut Parlar acknowledges the support from NSERC.

APPENDIX A. AUXILIARY RESULTS

Lemma 1. *For every $f, g \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$ and $0 \leq T, S \leq T_{\max}$, we have*

$$\left| \max_{0 \leq b \leq B} f(b, T) - \max_{0 \leq b \leq B} g(b, S) \right| \leq \max_{0 \leq b \leq B} |f(b, T) - g(b, S)| \quad \text{for every } B \in [0, B_{\max}].$$

Proof. Because f, g are continuous on the compact set $[0, B_{\max}] \times [0, T_{\max}]$, both maxima on the left are finite and attained at some $b_f, b_g \in [0, B]$, respectively. Therefore,

$$\begin{aligned} \max_{0 \leq b \leq B} f(b, T) - \max_{0 \leq b \leq B} g(b, S) &\leq f(b_f, T) - g(b_f, S) \leq \max_{0 \leq b \leq B} |f(b, T) - g(b, S)|, \\ \max_{0 \leq b \leq B} g(b, S) - \max_{0 \leq b \leq B} f(b, T) &\leq g(b_g, T) - f(b_g, S) \leq \max_{0 \leq b \leq B} |g(b, T) - f(b, S)|, \end{aligned}$$

which together complete the proof. \square

Lemma 2. *For every $f \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$, the function $F(B, T) := \max_{0 \leq b \leq B} f(b, T)$ is also in $\mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$.*

Proof. Because $f(\cdot, \cdot)$ is continuous on $[0, B_{\max}] \times [0, T_{\max}]$, it is uniformly continuous. For every $\varepsilon > 0$ there exists some $\delta > 0$ such that, whenever $(B, T), (C, S) \in [0, B_{\max}] \times [0, T_{\max}]$ satisfies $|B - C| + |T - S| \leq \delta$, we have $|f(B, T) - f(C, S)| \leq \varepsilon$. Note that

$$|F(B, T) - F(C, S)| \leq |F(B, T) - F(B, S)| + |F(B, S) - F(C, S)|.$$

By Lemma 1 with $g \equiv f$, we have $|F(B, T) - F(B, S)| \leq \max_{0 \leq b \leq B} |f(b, T) - f(b, S)| \leq \varepsilon$. We shall next show that $|F(B, S) - F(C, S)| \leq \varepsilon$ as well. Let us assume without loss of generality that $B \leq C$. Because $b \mapsto f(b, S)$ is continuous on the compact interval $[0, C]$, the value $F(C, S) = \max_{0 \leq b \leq C} f(b, S)$ is attained at some $b_C \in [0, C]$. If $b_C \in [0, B]$, then $|F(B, S) - F(C, S)| = 0 \leq \varepsilon$. If $b_C \notin [0, B]$, then we must have $b_C \in (B, C]$ and

$$0 \leq |F(B, S) - F(C, S)| = F(C, S) - F(B, S) \leq f(b_C, S) - f(B, S) \leq \varepsilon,$$

because $|b_C - B| = b_C - B \leq C - B \leq \delta$. \square

Lemma 3. *Let $f \in \mathcal{C}([0, B_{\max}]^2 \times [0, T_{\max}])$. Define*

$$F(B, T) = \max_{0 \leq b \leq B_{\max}} f(b, B, T) \quad \text{and} \quad I(b, B, T) = \int_0^T \lambda e^{-\lambda t} f(b, B, T - t) dt.$$

Then $F \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$ and $I \in \mathcal{C}([0, B_{\max}]^2 \times [0, T_{\max}])$.

Proof. For every $(B, T), (C, S) \in [0, B_{\max}] \times [0, T_{\max}]$, the maxima $F(B, T)$ and $F(C, S)$ are attained at some b_{BT} and b_{CS} in $[0, B_{\max}]$. Then

$$\begin{aligned} F(B, T) - F(C, S) &\leq f(b_{BT}, B, T) - f(b_{BT}, C, S) \leq \max_{0 \leq b \leq B_{\max}} |f(b, B, T) - f(b, C, S)|, \\ F(C, S) - F(B, T) &\leq f(b_{CS}, B, T) - f(b_{CS}, C, S) \leq \max_{0 \leq b \leq B_{\max}} |f(b, C, S) - f(b, B, T)|. \end{aligned}$$

Hence, $|F(B, T) - F(C, S)| \leq \max_{0 \leq b \leq B_{\max}} |f(b, B, T) - f(b, C, S)|$. Because f is uniformly continuous on the compact set $[0, B_{\max}]^2 \times [0, T_{\max}]$, one can find for every $\varepsilon > 0$ a number $\delta > 0$ such that $|b - c| + |B - C| + |T - S| \leq \delta$ implies $|f(b, B, T) - f(c, C, S)| \leq \varepsilon$; therefore, $|F(B, T) - F(C, S)| \leq \max_{0 \leq b \leq B_{\max}} |f(b, B, T) - f(b, C, S)| \leq \varepsilon$, and $F \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$. Without loss of generality, we can assume that $0 < \delta < \varepsilon$

(otherwise, replace δ with $\delta \wedge \varepsilon$). Then

$$\begin{aligned} & \left| \int_0^T \lambda e^{-\lambda t} f(b, B, T-t) dt - \int_0^S \lambda e^{-\lambda t} f(c, C, S-t) dt \right| \leq \int_0^T \lambda e^{-\lambda t} |f(b, B, T-t) - f(b, B, S-t)| dt \\ & + \left| \int_0^T \lambda e^{-\lambda t} f(b, B, S-t) dt - \int_0^S \lambda e^{-\lambda t} f(b, B, S-t) dt \right| + \int_0^S \lambda e^{-\lambda t} |f(b, B, S-t) - f(c, C, S-t)| dt \\ & \leq \varepsilon \int_0^T \lambda e^{-\lambda t} dt + \int_{T \wedge S}^{T \vee S} \lambda e^{-\lambda t} |f(b, B, S-t)| dt + \varepsilon \int_0^S \lambda e^{-\lambda t} dt \leq 2\varepsilon + \lambda \|f\| |T-S| \leq (2 + \lambda \|f\|) \varepsilon, \end{aligned}$$

which proves that $I \in \mathcal{C}([0, B_{\max}]^2 \times [0, T_{\max}])$. \square

Lemma 4. *If $w \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$, then $M_1 w \in \mathcal{C}([0, B_{\max}]^2 \times [0, T_{\max}])$ and $Mw \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$.*

Proof. If $w \in \mathcal{C}([0, B_{\max}] \times [0, T_{\max}])$, then both $(b, B, T) \mapsto w(B, T)$ and $(b, B, T) \mapsto w((B-b)^+, T) - w(B, T)$ are in $\mathcal{C}([0, B_{\max}]^2 \times [0, T_{\max}])$. The second part of Lemma 3 implies that both integrals in (5) are in $\mathcal{C}([0, B_{\max}]^2 \times [0, T_{\max}])$. Because the product and sum of continuous functions are continuous, the continuity of $M_1 w$ immediately follows. Finally, $(Mw)(B, T) = \max_{0 \leq b \leq B_{\max}} (M_1 w)(b, B, T)$ in (6) is continuous because of the first part of Lemma 3. \square

REFERENCES

- [1] S. Cholette, Ö. Özlük, and M. Parlar. Optimal keyword bids in search-based advertising with stochastic ad positions. *Journal of Optimization Theory and Applications*, 2011.
- [2] Emarketer.com. US online advertising spending 2008-2014, 2010. Available from http://www.emarketer.com/Reports/All/Emarketer_2000693.aspx. Accessed March 29, 2011.
- [3] G. E. Fruchter and W. Dou. Optimal budget allocation over time for keyword ads in web portals. *Journal of Optimization Theory and Applications*, 124(1):157–174, 2005.
- [4] Google. What is the adwords “quality score” and how is it calculated? Available at <http://adwords.google.com/support/aw/bin/answer.py?hl=en&answer=10215>. Accessed on September 22, 2011.
- [5] B. Kitts and B. Leblanc. Optimal Bidding on Keyword Auctions. *Electronic Markets*, 14(3):186–201, 2004.
- [6] Ö. Özlük. Search engine advertising: An overview from revenue management angle. Manuscript in progress, 2011.
- [7] Ö. Özlük and S. Cholette. Allocating expenditures across keywords in search advertising. *Journal of Revenue and Pricing Management*, 6(4):347–356, 2007.
- [8] P. Rusmevichientong and D. P. Williamson. An adaptive algorithm for selecting profitable keywords for search-based advertising services. In *Proceedings of the 7th ACM conference on Electronic commerce*, EC '06, pages 260–269, New York, NY, USA, 2006. ACM.
- [9] D. Szetela, J. Kerschbaum, and M. Flores. *Pay-Per-Click Search Engine Marketing: An Hour a Day*. Serious skills. John Wiley & Sons, 2009.

(S. Dayanik) BILKENT UNIVERSITY, DEPARTMENTS OF INDUSTRIAL ENGINEERING AND MATHEMATICS, BILKENT 06800, ANKARA, TURKEY (sdayanik@bilkent.edu.tr)

(M. Parlar) DEGROOTE SCHOOL OF BUSINESS, MCMASTER UNIVERSITY, HAMILTON, ONTARIO L8S 4M4, CANADA (parlar@mcmaster.ca)