ORIGINAL ARTICLE

Optimal dynamic multi-keyword bidding policy of an advertiser in search-based advertising

Savas Dayanik[1] · **Semih O. Sezer**[2]

© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2022

Abstract

Sponsored search advertisement allows advertisers to target their messages to appropriate customer segments at low costs. While search engines are interested in auction mechanisms that boost their revenues, advertisers seek optimal bidding strategies to increase their net sale revenues for multiple keywords under strict daily budget constraints in an environment where keyword query arrivals, competitor bid amounts, and user purchases are random. We focus on the advertiser's question and formulate her optimal intraday dynamic multi-keyword bidding problem as a continuous-time stochastic optimization problem. We solve the problem, characterize an optimal policy, and bring a numerical algorithm for implementation. We also illustrate our optimal bidding policy and its benefits over heuristic solutions on numerical examples.

1 Introduction

Sponsored search advertisement is a service of many internet search engines that allows businesses to communicate their messages to their potential customers along with the user search results. It is a popular marketing tool for advertisers as the target audience can be reached effectively at low costs. Parallel to its increasing popularity,

✉ Semih O. Sezer
sezer@sabanciuniv.edu

Savas Dayanik
sdayanik@bilkent.edu.tr

1 Industrial Engineering Department, Bilkent University, 06800 Bilkent, Ankara, Turkey

2 Faculty of Engineering and Natural Sciences, Sabancı University, 34956 Tuzla, Istanbul, Turkey

🖄 Springer

there has been a growing literature on search-based advertisement models, for which there are three different major streams of research. One stream brings a game theoretic perspective and studies the equilibria of the advertisers under various auction settings. The reader may refer to the review papers Maillé et al. (2012) and Qin et al. (2015), and also some of the recent papers Balseiro et al. (2015), Hummel (2018), Bae and Kagel (2019), Kotowski (2020) with their references for this line of work. A second stream of research focuses on the mechanism design problem of the search engine company with the objective of finding ad allocation/assignment policies to increase search engine's revenues. This line of work studies the problem as a general online resource allocation or matching problem. Earlier papers with such a focus include Mehta et al. (2007), Abrams et al. (2007), Goel and Mehta (2008), Devanur and Hayes (2009), Feldman et al. (2010), Goel et al. (2010), Mahdian et al. (2012). For recent results and discussions, the reader may refer to Jaillet and Lu (2014), Naor and Wajc (2015), Brubach et al. (2016), Devanur et al. (2019), Boulatov and Severinov (2021) and their references. A relatively less developed third and final stream of research considers the problem from the advertiser's perspective and aims at finding bidding policies maximizing the number of impressions, clicks, or revenues. This is the line of work we contribute to in the current paper.

In practice, an advertiser first compiles a list of keywords that she thinks are related to her products and services. Every time a search engine user queries a keyword, the search engine runs an online auction among the advertisers who keep the same keyword on their lists. After the bids are received sealed (namely, competitors do not see each others' bids), they are sorted in decreasing order according to a sorting algorithm taking into account the bid amounts and also a number of other factors like the relevancies of the ads, landing pages' contents, etc. Then the (chosen) sponsored links are displayed in the same order on the result page together with the organic search results; see, for example, Jansen and Mullen (2008) and Özlük (2011) for more details. The advertiser is charged only if the link is actually clicked, and the amount is calculated by the auction mechanism of the search engine (e.g., the generalized second price rule) with which, the actual cost is often less than the bid amount. Also, each advertiser fixes a finite daily budget as a protection against excessive spending, and the sponsored links cease to appear on the result pages as soon as the budget is depleted.

In the first line of research described above (with a game theoretic approach) the bidding behavior is generally governed by the equilibrium conditions, and in the second stream (taking search engine's perspective) bid amounts/policies are often assumed to be given. In real life, an advertiser typically acts as a revenue maximizing agent. Her problem of determining an optimal bidding policy is quite complicated and includes many stochastic elements: not only the keyword queries arrive randomly in continuous time, but also the number of competitors are unknown, and their bid amounts are unobserved. Hence, the positions of sponsored links on the search pages and the click events are all random to the advertiser. Furthermore, if the advertiser's link is displayed and clicked, the sales events and the sale amounts are also uncertain. Therefore, the advertiser faces a difficult optimization problem in a highly stochastic environment. The reader may refer to Pin and Key (2011) for further details and comments on the randomness in an ad auction environment.

In the current paper, we consider a stochastic model for this environment. Extending the single keyword formulation of Dayanik and Parlar (2013), we study the problem of finding a bidding policy for multiple keywords in a general auction setting as a continuous-time optimization problem, and we solve it for an optimal policy using stochastic dynamic programming. This is the main contribution of the paper.

Earlier work (taking the advertiser's perspective) include Kitts and Leblanc (2004) determining the optimal bid amounts for multiple keywords in an open bid system by solving a deterministic integer program (with unknown functions estimated). In a similar setting with known competitors' bids, a different formulation to determine the number of times the advertiser participates in different auctions is considered in Chaitanya and Narahari (2012). Today, however, an overwhelming fraction of the sponsored search auctions are with closed bids.

Özlük and Cholette (2007), Cholette et al. (2012), Abhishek and Hosanagar (2013), and Küçükaydın et al. (2020) study the bidding problem in closed bid generalized second price (GSP) auctions. These papers determine a constant bid price for the whole day for each keyword under the *soft* budget constraint; namely, the *expected* total daily spending should not exceed the available budget. A simple randomized bidding policy is considered in Feldman et al. (2007), again under a soft budget constraint. Taking a different objective, Selçuk and Özlük (2013) minimizes the expected spending while satisfying a certain level of exposure. Although providing important insights and contributions, in these models with soft constraints, the budget limits may be exceeded with non-negligible probabilities, and this can be a problem in practice. Furthermore, constant bids or more general time-stationary bidding strategies may cause profit losses because they do not adjust themselves as the keyword arrival, remaining time, and budget processes are continuously observed along the day.

For a single keyword, the first optimal dynamic bidding strategy in a stochastic model with closed bids and under strict budget constraint was derived by the aforementioned paper Dayanik and Parlar (2013). The straightforward generalization of that model to K-many keywords with a given budget B over a planning horizon T would require solving the problem

$$\sup_{B_1,\ldots,B_K} \sum_{k=1}^{K} f_k(B_k, T) \qquad \text{subject to the constraint} \quad \sum_{k=1}^{K} B_k \leq B,$$

where the function f_k denotes the expected revenue from the keyword k that one could obtain using the solution of the single keyword formulation. That is, we first solve K-many two-dimensional problems to obtain the functions f_ks, and then, as the next step, we determine the optimal partitioning of the budget to different keywords (which can, for example, be obtained using a dynamic programming approach over the keywords). However, such an approach not only suffers from the increased computation times as the number of keywords grows, but also has the drawback that some of the budget allocated for certain keywords may remain idle as a result of low query traffic realizations (and not used for other keywords with high volumes).

Here, in this paper, we show that, independent of the number of keywords, the problem can be solved using a two-dimensional budget-time jump process. We char-

acterize the value function as the unique solution of a differential equation on this two-dimensional space and show how to numerically obtain it; see (34). This approach shields us from the computational issues and suboptimality of the straightforward extension of Dayanik and Parlar (2013) described above. To highlight the importance of using such an approach, in Sect. 6, we numerically compare the performance of our multi-keyword optimal policy to a number of heuristics in which the budget is split for each keyword and one-keyword policy of Dayanik and Parlar (2013) is applied separately. These heuristics differ in the way they partition the budget for different keywords. In our results, we observe that the multi-keyword optimal policy significantly outperforms the heuristic ones when the available advertisement budget is small or moderate. As expected, when the budget gets large, this difference becomes less significant. In our setup, to close up the expected net revenues between optimal and heuristic bidding policies, one needs to almost double the optimal daily advertisement budget.

The current paper differs from Dayanik and Parlar (2013) on other accounts as well. Firstly, we consider a general pricing mechanism under which the advertiser pays less than her bid amount (which was overlooked in the cited paper). Secondly, we consider a general *display-and-click* probability function for each keyword rather than assuming a special distribution for the location of the ad on the result page and then expressing the conditional click probability dependent on this location. Using display-and-click probability functions gives us a direct and flexible way of modeling the effects of bids. These functions can also be regarded as determined by a game model, in which the dynamics of the game determine the click probabilities for a given bid amount. As the third difference of the current paper, we use a different dynamic programming operator which does not assume that a keyword query is realized at time zero and therefore reflects the true nature of the problem better. Overall, the current paper greatly improves the analysis and formulation of Dayanik and Parlar (2013) and provides a more general picture of this continuous-time optimization problem in its multi-keyword setting.

As in many decision problems, applying the results of the current paper in a real life scenario requires reliable estimates of the parameters and distributions. When an advertiser launches a new campaign, these estimates may not be readily available. To address such issues, the current model should be enlarged to include a learning window in the beginning of the campaign over which the advertiser places bids and enters the auctions. The study of such an extended formulation is clearly non-trivial and left as future research. For related work and discussions, the reader may refer to Rusmevichientong and Williamson (2006), Borgs et al. (2007), Perlich et al. (2012), Skiera and Nabout (2013), Iyer et al. (2014), Lee et al. (2017), and the references therein.

The rest of the paper is organized as follows. In Sect. 2, we present a stochastic model for the bidding process and formulate the optimal bidding problem. The dynamic programming operator associated with this formulation is discussed next in Sect. 3 along with its useful properties. In Sect. 4, we construct sequential approximations for the value function, and we present an optimal bidding policy for the advertiser. For practical implementations, successive approximations turn out to be a computationally inefficient way to calculate the value function. Instead, in Sect. 5,

we derive a differential equation that is much faster to solve for the value function. Finally, in Sect. 6, we illustrate on some numerical results the computation of maximum expected total net daily revenues and optimal bidding policy, and we quantify their benefits over the simpler heuristics. Appendix at the end includes supplementary proofs.

2 Problem formulation

Suppose that an advertiser compiles a fixed list of K keywords. Let $\tau_1, \tau_2, \ldots$ be the successive query times and $\kappa_1, \kappa_2, \ldots$ be the keywords searched at those times, respectively. The collection $(\tau_n, \kappa_n)_{n \geq 1}$ forms a Poisson point process

$$p((0, t] \times \{k\}) = \sum_{n \geq 1} 1_{\{\tau_n \leq t, \, \kappa_n = k\}}, \quad \text{for } t \geq 0 \quad \text{and} \quad 1 \leq k \leq K,$$

with compensator $\mathbb{E}[p((0, t] \times \{k\})] = \lambda t \times \frac{\lambda_k}{\lambda}$, for $1 \leq k \leq K$, where $\lambda = \sum_{k=1}^{K} \lambda_k$, and with filtration denoted by $\mathbb{F}^p = \{\mathcal{F}_t^p\}_{t \geq 0}$ in the sequel. In plain words, queries arrive according to a simple Poisson process with intensity λ, and keywords have the distributions $\mathbb{P}(\kappa_n = k) = \lambda_k / \lambda$, $1 \leq k \leq K$, independently for all $n \geq 1$.

The advertiser starts with an initial budget $0 \leq B \leq B_{\max}$ allocated for a period $[0, T]$ for some $0 \leq T \leq T_{\max}$, and she bids for an ad position on the result page after every query. We assume that the advertiser is charged only when the ad is clicked. The amount charged is less than the bid amout and determined according to a given auction pricing rule. More precisely, we let b_n be the bid amount at the n'th query and $b_n R_n$ is the actual price paid upon click event, where $R_n \in (0, 1]$ denotes the effect of the pricing mechanism (for example, in a second price auction, it is the ratio of the largest bid smaller than b_n to b_n). Also, we let Z_n be the indicator of the event that the ad is displayed and clicked at the n'th query; that is, $Z_n = 1$ if the ad is displayed on the result page and clicked (i.e., a search user visits the advertiser's website), and it is zero otherwise. On the click event, the user potentially generates a revenue, and the amount is denoted by the random variable W_n.

On this environment, as the advertiser places bids, she also observes the click events, her actual payments, and the revenue generated by the users. All these observations are represented by the filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t \geq 0}$, which is obtained by augmenting $\mathbb{F}^p$ with previously observed random variables $(Z_., Z.R., Z.W.)$'s. For notational convenience, we omit the dependence of the observation filtration on the bidding policy.

At a given query time τ_n, $n \geq 1$, given the observations $\mathcal{F}_{\tau_n}$, the conditional distribution of R_n depends on κ_n only, and the conditional distributions of Z_n depends on κ_n and b_n only. That is, in the $\mathbb{P}$-almost sure sense, we have

$$\mathbb{P}[R_n \leq r \mid \mathcal{F}_{\tau_n}] = H_{\kappa_n}(r) \quad \text{and} \quad \mathbb{P}[Z_n = 1 \mid \mathcal{F}_{\tau_n}] = G_{\kappa_n}(b_n) \tag{1}$$

in terms of the given functions $r \mapsto H_k(r)$ and $b \mapsto G_k(b)$, for $1 \leq k \leq K$. For the keyword k, the function H_k gives us the (conditional) cumulative distribution of the

discount random variable R_n, and G_k is the (conditional) display-and-click probability function for the binary random variable Z_n. We additionally assume that R_n and Z_n are conditionally independent given $\mathcal{F}_{\tau_n}$.

We denote

$$\rho_k := \int_0^1 r\, H_k(\mathrm{d}r), \qquad \text{for every } 1 \le k \le K. \tag{2}$$

For the display-and-click probability functions G_k's, for $k \le K$, we take the natural assumptions that i) $G_k(0) = 0$, ii) $G_k(b)$ is non-decreasing in b, and ii) small increments in the bid amounts do not drastically increase the click probabilities. That is, for each keyword, the function $b \mapsto G_k(b)$ is a non-decreasing continuous function starting from the origin.

Given $\mathcal{F}_{\tau_n}$, independent from the random variables R_n and Z_n, the potential random revenue W_n has a conditional mean

$$\mathbb{E}[W_n \mid \mathcal{F}_n] = \sum_{k=1}^K \mu_k \mathbf{1}_{\{\kappa_n = k\}} = \mu_{\kappa_n}, \qquad \mathbb{P}\text{-almost surely}, \tag{3}$$

where μ_k denotes the expected revenue generated by the user searching the keyword k and clicking the advertiser's ad. We have $\mu_k \in (0, \infty)$ for all $1 \le k \le K$, and when needed, we denote their maximum as $\bar{\mu} := \max_k \mu_k$.

In this setup, for every

$$(B, T) \in \Delta := [0, B_{\max}] \times [0, T_{\max}], \tag{4}$$

the objective of the advertiser is to compute the maximum expected net revenue

$$V(B, T) := \sup_{(b_n)_{n \ge 1} \in \mathscr{D}(B,T)} \mathbb{E}\left[\sum_{n \ge 1} (W_n - b_n R_n) Z_n \mathbf{1}_{\{\tau_i \le T\}} \right], \tag{5}$$

in which the supremum is taken over the collection of *admissible* bidding strategies

$$\mathscr{D}(B, T) := \left\{ (b_n)_{n \ge 1}; \ b_n \in \mathcal{F}_{\tau_n}, \ b_n \ge 0 \text{ for every } n \ge 1, \text{ and} \right.$$

$$\left. \sum_{n=1}^\infty b_n R_n Z_n \mathbf{1}_{\{\tau_n \le T\}} \le B \right\} \tag{6}$$

imposing the strict budget constraint and also the measurability of the bids (so that looking into future is not allowed).

In the appendix, we show that

$$\mathbb{E}\left[\sum_{n \ge 1} (W_n - b_n R_n) Z_n \mathbf{1}_{\{\tau_n \le T\}} \right] = \mathbb{E}\left[\sum_{n \ge 1} (\mu_{\kappa_n} - b_n \rho_{\kappa_n}) G_{\kappa_n}(b_n) \mathbf{1}_{\{\tau_n \le T\}} \right], \tag{7}$$

which gives us an alternative representation for the expected net revenue associated with a given bidding policy.

Remark 1 The representation in (7) implies that it is never optimal to bid more than μ_k/ρ_k when the keyword observed is k. Therefore, the supremum in (5) can be taken over admissible policies for which $b_n \leq \mu_{\kappa_n}/\rho_{\kappa_n}$ for all $n \geq 1$.

Remark 2 For $B_1 \leq B_2 \leq B_{\max}$, we have $\mathscr{D}(B_1, T) \subseteq \mathscr{D}(B_2, T)$. Hence the mapping $B \mapsto V(B, T)$ is non-decreasing. Similarly, for $T_1 \leq T_2 \leq T_{\max}$, $\mathscr{D}(B, T_1) \subseteq \mathscr{D}(B, T_2)$ and the mapping $T \mapsto V(B, T)$ is non-decreasing as well.

Remark 3 The value function V is clearly non-negative because $b_n = 0$ for all $n \geq 1$ is an admissible policy for any pair $(B, T) \in \Delta$. Also, since $B \mapsto V(B, T)$ is non-decreasing, using the representation in (7) and the observation in Remark 1, we write

$$
\begin{aligned}
\lim_{B \to \infty} V(B, T) &= \sup_{B \geq 0} V(B, T) \\
&= \sup_{B \geq 0} \sup_{\substack{b_n \in \mathcal{F}_{\tau_n} \\ 0 \leq b_n \leq \frac{\mu_{\kappa_n}}{\rho_{\kappa_n}}}} \mathbb{E}\left[\left(\sum_{n \geq 1} (\mu_{\kappa_n} - b_n \rho_{\kappa_n}) G_{\kappa_n}(b_n) 1_{\{\tau_n \leq T\}}\right) \right. \\
&\quad \left. \times 1_{\{\sum_{n=1}^{\infty} b_n R_n Z_n 1_{\{\tau_n \leq T\}} \leq B\}}\right] \\
&= \sup_{\substack{b_n \in \mathcal{F}_{\tau_n} \\ 0 \leq b_n \leq \frac{\mu_{\kappa_n}}{\rho_{\kappa_n}}}} \sup_{B \geq 0} \mathbb{E}\left[\left(\sum_{n \geq 1} (\mu_{\kappa_n} - b_n \rho_{\kappa_n}) G_{\kappa_n}(b_n) 1_{\{\tau_n \leq T\}}\right) \right. \\
&\quad \left. \times 1_{\{\sum_{n=1}^{\infty} b_n R_n Z_n 1_{\{\tau_n \leq T\}} \leq B\}}\right] \\
&= \sup_{\substack{b_n \in \mathcal{F}_{\tau_n} \\ 0 \leq b_n \leq \frac{\mu_{\kappa_n}}{\rho_{\kappa_n}}}} \mathbb{E}\left[\sum_{n \geq 1} (\mu_{\kappa_n} - b_n \rho_{\kappa_n}) G_{\kappa_n}(b_n) 1_{\{\tau_n \leq T\}}\right] \equiv V(\infty, T).
\end{aligned}
$$

Note that the mapping $b \mapsto (\mu_k - b\rho_k) G_k(b)$ is continuous for every k, and therefore it attains its maximum on the interval $[0, \mu_k/\rho_k]$. If we let $b_{\infty,k}$ denote the smallest maximizer on this interval, it follows that it is optimal to bid b_{∞,κ_n} at time τ_n in the absence of any budget constraint, and this gives

$$
V(\infty, T) = \mathbb{E}\left[\sum_{n \geq 1} (\mu_{\kappa_n} - b_{\infty,\kappa_n} \rho_{\kappa_n}) G_{\kappa_n}(b_{\infty,\kappa_n}) 1_{\{\tau_n \leq T\}}\right] = \gamma \lambda T, \qquad (8)
$$

where $\gamma := \sum_{k=1}^{K} (\mu_k - b_{\infty,k} \rho_k) G_k(b_{\infty,k}) \frac{\lambda_k}{\lambda}$. Hence, we have the bounds $0 \leq V(B, T) \leq V(\infty, T) = \gamma \lambda T$.

Having a finite upper bound as described in the remark above is not surprising because, even if we could bid for keywords large enough to attract every web surfer who search them, the total number of web searches for those keywords is determined by the probability laws that do not depend on the advertisement campaign.

Remark 4 For a fixed budget $B \leq B_{\max}$, take two time horizons $0 \leq T_1 < T_2$ and let $b^{(\varepsilon)}$ be an ε-optimal policy for $V(B, T_2)$ for some $\varepsilon > 0$; that is,

$$V(B, T_2) \leq \mathbb{E}\left[\sum_{n \geq 1}(\mu_{\kappa_n} - b_n^{(\varepsilon)}\rho_{\kappa_n})G_{\kappa_n}(b_n^{(\varepsilon)})1_{\{\tau_n \leq T\}}\right] + \varepsilon.$$

The truncated bidding policy $\bar{b}^{(\varepsilon)}$ with $\bar{b}_n^{(\varepsilon)} = b_n^{(\varepsilon)}1_{\{\tau_n \leq T_1\}}$, $n \in \mathbb{N}$, is in $\mathscr{D}(B, T_1)$, and we have

$$0 \leq V(B, T_2) - V(B, T_1) \leq \mathbb{E}\left[\sum_{n \geq 1}(\mu_{\kappa_n} - b_n^{(\varepsilon)}\rho_{\kappa_n})G_{\kappa_n}(b_n^{(\varepsilon)})1_{\{T_1 < \tau_n \leq T_2\}}\right] + \varepsilon$$

$$\leq \mathbb{E}\left[\sum_{n \geq 1}\mu_{\kappa_n}1_{\{T_1 < \tau_n \leq T_2\}}\right] + \varepsilon \leq \bar{\mu}\,\mathbb{E}[N_{T_2} - N_{T_1}] + \varepsilon = \bar{\mu}\lambda(T_2 - T_1) + \varepsilon,$$

where N denotes the counting process $N_t = \sum_{n \geq 1}1_{\{\tau_n \leq T\}}$ for $t \geq 0$. Since ε above is arbitrary, it follows that the function $T \mapsto V(B, T)$ is Lipschitz continuous.

3 Dynamic programming operator

Let us introduce the budget process $\{B_t\}_{t \in [0,T]}$, where B_t represents the remaining budget at time t. Clearly, for a given bidding policy $(b_n)_{n \geq 1}$, we have

$$B_t = \begin{cases} B_{\tau_n}, & \text{for } t \in [\tau_n, \tau_{n+1}), \\ B_{\tau_{n+1}-} - Z_{n+1}R_{n+1}b_{n+1}, & \text{for } t = \tau_{n+1}, \end{cases} \tag{9}$$

for $n \geq 0$, with $B_0 = B$ and $\tau_0 = 0$. The principle of dynamic programming suggests that if the first keyword search occurs before time T with a bid amount b_1, then the realized net profit is $(W_1 - b_1R_1)Z_1$ and the optimal conditional expected revenue collected from then on should be given by $V(B_{\tau_1}, T - \tau_1) = V(B - b_1R_1Z_1, T - \tau_1)$. Therefore, the value function should satisfy the equation $V(B, T) = \mathcal{D}[V](B, T)$ in terms of the operator

$$D[f](B, T) = \sup_{\substack{b_1 \in \mathcal{F}_{\tau_1} \\ 0 \leq b_1 \leq B}} \mathbb{E}\left[1_{\{\tau_1 \leq T\}}\Big((W_1 - b_1R_1)Z_1 + f(B - b_1R_1Z_1, T - \tau_1)\Big)\right] \tag{10}$$

defined for Borel functions f's on Δ.

Proposition 1 *If f_1 and f_2 are two functions for which $0 \leq f_1(\cdot, T) \leq f_2(\cdot, T) \leq \gamma\lambda T$ (see Remark 3 for the definition of γ), then we have $0 \leq \mathcal{D}[f_1](\cdot, T) \leq \mathcal{D}[f_2](\cdot, T) \leq \gamma\lambda T$.*

As $\tau_1, \kappa_1, b_1 \in \mathcal{F}_{\tau_1}$, conditioning on $\mathcal{F}_{\tau_1}$ yields

$$
\mathbb{E}\left[1_{\{\tau_1 \leq T\}} f(B - Z_1 R_1 b_1, T - \tau_1)\right] = \mathbb{E}\left[1_{\{\tau_1 \leq T\}}\left(f(B, T - \tau_1) + G_{\kappa_1}(b_1)\right.\right.
$$
$$
\left.\left.\times \left(\int_0^1 f(B - rb_1, T - \tau_1) H_{\kappa_1}(dr) - f(B, T - \tau_1)\right)\right)\right],
$$

which implies that the operator $\mathcal{D}$ in (10) can be rewritten as

$$
\mathcal{D}[f](B, T) = \sup_{b_1 \in \mathcal{F}_{\tau_1}} \mathbb{E}\left[1_{\{\tau_1 \leq T\}}\left(f(B, T - \tau_1) + G_{\kappa_1}(b_1)\right.\right.
$$
$$
\left.\left.\times \left(\mu_{\kappa_1} - b_1 \rho_{\kappa_1} + \int_0^1 f(B - rb_1, T - \tau_1) H_{\kappa_1}(dr) - f(B, T - \tau_1)\right)\right)\right]. \tag{11}
$$

For notational convenience, let us introduce

$$
\Gamma_k[f](b, B, s) := \mu_k - b\rho_k + \int_0^1 f(B - rb, s) H_k(dr) - f(B, s)
$$
$$
M_k[f](b, B, s) := G_k(b) \cdot \Gamma_k[f](b, B, s) \tag{12}
$$

and define

$$
M_k^*[f](B, s) := \sup_{b \leq B} M_k[f](b, B, s), \quad \text{for } 1 \leq k \leq K \text{ and } (B, s) \in \Delta. \tag{13}
$$

The operator Γ_k represents the incremental contribution/reward of the keyword k, on the display-and-click event, under the bid amount b when the value of continuing is given by the function f. M_k takes the product of this incremental reward with the probability $G_k(b)$ of the display-and-click event, and M_k^* seeks for the value of b maximizing this product. Note that, for a function f for which $B \mapsto f(B, T)$ is non-decreasing, Γ_k is non-increasing in the bid amount as expected (higher bid reduces the budget more if the ad is clicked). However, the trade-off comes from the property of G_k that it is non-decreasing in b. Hence, solving for the best b in (13) a non-trivial one-dimensional optimization problem.

In (12–13), if we replace f with the value function V in (5) we expect the maximizer of $M_k^*[V](B, s)$, if exists, be the optimal bid amount if the keyword k is searched when there is s units of time left (until the end of time horizon) with the available budget being B. We verify this intuition rigorously later in Sect. 4. In this section, f denotes an arbitrary function with which we continue after the first query in the one-step operator $\mathcal{D}$.

If $f(\cdot, \cdot)$ is a continuous function on Δ, the proof of Proposition 1 (i–ii) in Dayanik and Parlar (2013) can be modified easily to prove that the supremum in (13) is attained, and $M_k^*[f](\cdot, \cdot)$ is continuous on Δ for each $1 \leq k \leq K$. In this case, in terms of

$$
b_k^*[f](B, s) := \arg\max_{b \in [0, B]} M_k[f](b, B, s), \tag{14}
$$

the supremum in (11) is attained when we set $b_1 = b_{\kappa_1}^*[f](B, T - \tau_1)1_{\{\tau_1 \leq T\}} \in \mathcal{F}_{\tau_1}$, and we have the explicit form

$$
\begin{aligned}
\mathcal{D}[f](B, T) &= \mathbb{E}\left[1_{\{\tau_1 \leq T\}}\Big(f(B, T - \tau_1) + M_{\kappa_1}^*[f](B, T - \tau_1)\Big)\right] \\
&= \sum_{k=1}^{K} e^{-\lambda T} \int_0^T \lambda_k e^{\lambda u} \Big[f(B, u) + M_k^*[f](B, u)\Big] \mathrm{d}u.
\end{aligned}
\tag{15}
$$

Corollary 1 *The explicit characterization in* (15) *shows that* $\mathcal{D}[f]$ *is continuous on* Δ *if so is* f.

In the sequel, we use $\mathcal{C}(\Delta)$ to denote the set of continuous functions on Δ. Proposition 2 strengthens the result in Corollary 1 by controlling the growth in each dimension under some assumptions. The proof the result is deferred to the appendix.

Proposition 2 1. *Suppose f is non-negative, bounded, and $T \mapsto f(B, T)$ is non-decreasing for every $B \in [0, B_{\max}]$. Then,*

$$
0 \leq \mathcal{D}[f](B, T_2) - \mathcal{D}[f](B, T_1) \leq \lambda(\bar{\mu} + \|f\|)(T_2 - T_1), \qquad \text{for } T_1 < T_2,
\tag{16}
$$

again for every $B \in [0, B_{\max}]$, where $\|f\| := \sup_{(B,T) \in \Delta} |f(B, T)|$.
2. *Assume that f is non-negative, bounded, and $B \mapsto f(B, T)$ is non-decreasing for every $T \in [0, T_{\max}]$. Also, suppose that there exist Lipschitz constants $\alpha_f > 0$ and $\alpha_G > 0$ for which*

$$
\begin{aligned}
f(B_2, T) - f(B_1, T) &\leq \alpha_f(B_2 - B_1), \qquad \text{for } B_1 < B_2 \text{ for every } T \in [0, T_{\max}], \\
\text{and } G_k(b_2) - G_k(b_1) &\leq \alpha_G(b_2 - b_1), \qquad \text{for } b_1 < b_2 \text{ for all } 1 \leq k \leq K.
\end{aligned}
$$

Then, for $B_1 < B_2$,

$$
0 \leq \mathcal{D}[f](B_2, T) - \mathcal{D}[f](B_1, T) \leq (B_2 - B_1)(1 - e^{-\lambda T_{\max}})\big[\alpha_f + \alpha_G(\bar{\mu} + \|f\|)\big],
\tag{17}
$$

for every $T \in [0, T_{\max}]$.

Lemma 1 below is analogous to Proposition 1 (iv) in Dayanik and Parlar (2013). Its proof is similar, hence omitted.

Lemma 1 *The dynamic programming operator D is a contraction mapping and we have $\|\mathcal{D}[f_1] - \mathcal{D}[f_2]\| \leq (1 - e^{-\lambda T_{max}})\|f_1 - f_2\|$ for two functions $f_1, f_2 \in \mathcal{C}(\Delta)$. Therefore, if $\mathcal{D}$ has a fixed point in $C(\Delta)$, then it must be unique.*

If $B \mapsto f(B, T)$ is continuous, bounded, and non-decreasing, then $b \mapsto \Gamma_k[f](b, B, T)$ is continuous (thanks to the bounded convergence theorem) and strictly

decreasing. Also, we have $\Gamma_k[f](0, B, T) = \mu_k > 0$. Hence, the region of bids with non-negative contributions

$$A[f](B, T) := \{b \in [0, B]; \; \Gamma_k[f](b, B, T) > 0\} \tag{18}$$

is a non-empty interval starting from the origin having a strictly positive (open) right boundary. We have

$$\sup_{b \leq B} M_k[f](b, B, T) = \sup_{b \in A[f](B,T)} M_k[f](b, B, T), \tag{19}$$

and, unless the click probability $G_k(b)$ is zero for all $b \in A[f](B, T)$, the maximizer $b_k^*[f](B, s)$ in (14) is strictly positive; that is, 'no bidding' ($b = 0$) is a suboptimal choice.

Note that when $B \mapsto f(B, T)$ is non-decreasing,

$$\int_0^1 f(B - rb, s) \, H_k(dr) - f(B, s) = \int_0^1 \left[f(B - rb, s) - f(B, s) \right] H_k(dr) \leq 0,$$

and therefore $\Gamma_k[f](b, B, T) \leq \mu_k - b\rho_k$. This further implies that $A[f](B, T) \subseteq [0, \frac{\mu_k}{\rho_k})$, which is in line with the upper bound on the bid amounts established in Remark 1.

4 Successive approximations and an optimal bidding policy

Let us construct the sequences of functions $(V_n)_{n\geq0}$, $(\underline{U}_n)_{n\geq0}$ and $(\overline{U}_n)_{n\geq0}$ on Δ such that

$$
\begin{aligned}
V_n(B, T) &:= \sup_{(b_i)_{i\geq1} \in \mathscr{D}(B,T)} \mathbb{E}\left[\sum_{i \leq n} (W_i - b_i R_i) Z_i 1_{\{\tau_i \leq T\}} \right], \\
\underline{U}_0(B, T) &:= 0 \quad \text{and} \quad \underline{U}_{n+1}(B, T) := \mathcal{D}[\underline{U}_n](B, T), \\
\overline{U}_0(B, T) &:= \gamma\lambda T \quad \text{and} \quad \overline{U}_{n+1}(B, T) := \mathcal{D}[\overline{U}_n](B, T), \quad \text{for } n \geq 0.
\end{aligned}
\tag{20}
$$

The first sequence is obtained by limiting the number bids, and the second and the third follow from the repeated applications of the dynamic programming operator $\mathcal{D}$ starting with a lower and upper bound on the original value function.

Remark 5 Note that $\underline{U}_0$ and $\overline{U}_0$ are both continuous on Δ. Because the operator $\mathcal{D}$ preserves the continuity (see Corollary 1), it follows by induction that the functions $\underline{U}_n$ and $\overline{U}_n$ are also continuous for each $n \geq 1$.

Proposition 3 *The sequences $(\underline{U}_n)_{n\geq0}$ and $(\overline{U}_n)_{n\geq0}$ are non-decreasing and non-increasing respectively. They converge to the same function $U_\infty = \lim_n \underline{U}_n = \lim_n \overline{U}_n$ uniformly on Δ with the error bounds*

$$\|U_\infty - \overline{U}_n\| \leq \gamma\lambda T_{max} \, e^{\lambda T_{max}} \left(1 - e^{-\lambda T_{max}}\right)^n$$

$$and \quad \|U_\infty - \underline{U}_n\| \leq \gamma \lambda T_{max} e^{\lambda T_{max}} \left(1 - e^{-\lambda T_{max}}\right)^n, \tag{21}$$

for every $n \geq 0$. The limit function U_∞ is the unique fixed point of the operator $\mathcal{D}$ on $\mathcal{C}(\Delta)$.

Remark 6 Both $\underline{U}_0$ and $\overline{U}_0$ satisfy the assumptions in Proposition 2 (i) and we have the bounds $0 \leq \underline{U}_n \leq \overline{U}_n \leq \gamma \lambda T_{max}$ for all $n \geq 0$. Hence, it follows by induction that, for $T_1 < T_2$ and $B \in [0, B_{max}]$,

$$\begin{aligned} 0 \leq \overline{U}_n(B, T_2) - \overline{U}_n(B, T_1) &\leq \lambda(\bar{\mu} + \gamma \lambda T_{max})(T_2 - T_1) \\ 0 \leq \underline{U}_n(B, T_2) - \underline{U}_n(B, T_1) &\leq \lambda(\bar{\mu} + \gamma \lambda T_{max})(T_2 - T_1) \end{aligned} \tag{22}$$

for all $n \geq 1$. Letting $n \to \infty$ in (22) shows that the same bounds also hold for U_∞.

Similarly, because $\underline{U}_0$ and $\overline{U}_0$ satisfy the assumptions in Proposition 2 (ii) with the common upper bound $\gamma \lambda T_{max}$, again it can be shown inductively that, when the click-and-display probability functions are all Lipschitz continuous with a Lipschitz constant α_G, we have for $B_1 < B_2$ and $T \in [0, T_{max}]$

$$\begin{aligned} 0 \leq \overline{U}_n(B_2, T) - \overline{U}_n(B_1, T) &\leq (B_2 - B_1)(1 - e^{-\lambda T_{max}})\alpha_G(\bar{\mu} + \gamma \lambda T_{max}) \\ &\times \sum_{i=0}^{n-1}(1 - e^{-\lambda T_{max}})^i \\ 0 \leq \overline{U}_n(B_2, T) - \overline{U}_n(B_1, T) &\leq (B_2 - B_1)(1 - e^{-\lambda T_{max}})\alpha_G(\bar{\mu} + \gamma \lambda T_{max}) \\ &\times \sum_{i=0}^{n-1}(1 - e^{-\lambda T_{max}})^i, \end{aligned} \tag{23}$$

for all $n \geq 1$, and letting $n \to \infty$ above gives

$$0 \leq U_\infty(B_2, T) - U_\infty(B_1, T) \leq (B_2 - B_1)(e^{\lambda T_{max}} - 1)\alpha_G(\bar{\mu} + \gamma \lambda T_{max}). \tag{24}$$

Proposition 4 $V_n \nearrow V$ uniformly on Δ and we have

$$V(B, T) \geq V_n(B, T) \geq V(B, T) - \gamma \frac{(\lambda T_{max})^{n+1}}{(n+1)!}, \quad for \ n \geq 0. \tag{25}$$

Proof The first inequality in (25) is obvious. To establish the second, we note that a given feasible bidding policy $(b_i)_{i \geq 1}$ can be truncated after the n'th query and the expected revenue generated by the remaining bids are bounded from above by

$$E\Big[\sum_{i \geq n+1}(\mu_{\kappa_i} - b_{\infty,\kappa_i}\rho_i)G_{\kappa_i}(b_{\infty,\kappa_i})1_{\{\tau_i \leq T\}}\Big] = \mathbb{E}\Big[\sum_{i \geq n+1}\gamma 1_{\{\tau_i \leq T\}}\Big]$$
$$\leq \gamma \sum_{i \geq n+1}\mathbb{E}\big[1_{\{\tau_i \leq T_{max}\}}\big].$$

The i'th arrival time τ_i has the Erlang distribution. Using this distribution for the last term above, it is easy to show that $\sum_{i \geq n+1} \mathbb{E}[1_{\{\tau_i \leq T_{\max}\}}]$ is bounded above by $(\lambda T_{\max})^{n+1}/(n+1)!$. Hence, we have the following upper bound

$$E\Big[\sum_{i \geq 1}(W_i - b_i R_i)Z_i 1_{\{\tau_i \leq T\}}\Big] \leq E\Big[\sum_{i \leq n}(W_i - b_i R_i)Z_i 1_{\{\tau_i \leq T\}}\Big] + \gamma \frac{(\lambda T_{\max})^{n+1}}{(n+1)!}$$

for the expected revenue generated by a bidding policy $(b_i)_{i \geq 1}$. Taking next the supremum of both sides gives now the second inequality in (25). □

Lemma 2 *For any admissible policy $(b_i)_{i \geq 1}$, we have*

$$\mathbb{E}\Big[\sum_{i=1}^{n}(W_i - b_i R_i)Z_i 1_{\{\tau_i \leq T\}}\Big] \leq \mathbb{E}\Big[\sum_{i=1}^{n-j+1}(W_i - b_i R_i)Z_i 1_{\{\tau_i \leq T\}} + 1_{\{\tau_{n-j+1} \leq T\}}$$
$$\times \underline{U}_{j-1}(B_{\tau_{n-j+1}}T - \tau_{n-j+1})\Big], \qquad (26)$$

for all $1 \leq j \leq n+1$ (with $\tau_0 = 0$ for $j = n+1$).

For all $n \geq 1$ and $(B, T) \in \Delta$, let us define the bidding policy $b^{(n)} \equiv (b_i^{(n)})_{i \geq 1}$ with the corresponding budget policy $B^{(n)}$ (according to (9)) recursively as

$$b_i^{(n)} := \begin{cases} b_{\kappa_i}^*[\underline{U}_{n-i}](B_{\tau_{i-1}}^{(n)}, T - \tau_i) \cdot 1_{\{\tau_i \leq T\}}, & i \leq n, \\ 0, & i \geq n+1 \end{cases}, \qquad \text{for } i \geq 1, \qquad (27)$$

in terms of $b^*[\cdot](\cdot, \cdot)$ given in (14), and with $\tau_0 = 0$ for $i = 1$. Proposition 5 below shows that the we have $V_n = \underline{U}_n$ and the supremum in (20) is attained if we apply the bidding policy in (27). Clearly, both $b^{(n)}$ and $B^{(n)}$ depend on the initial point $(B, T) \in \Delta$. Here, we omit this dependence for notational convenience only.

Proposition 5 *For every $n \geq 1$ and $(B, T) \in \Delta$, we have*

$$V_n(B, T) = \underline{U}_n(B, T) = \mathbb{E}\left[\sum_{i=1}^{n}\Big(W_i - b_i^{(n)}R_i\Big)Z_i 1_{\{\tau_i \leq T\}}\right], \qquad (28)$$

where $(b_i^{(n)})_{i \geq 1}$ is defined in (27).

The results in Lemma 2 and Proposition 5 are indeed intuitive from a dynamic programming point of view. Lemma 2 shows (with the help of Proposition 5) that any feasible n-bid policy can be improved after any bid by switching to the optimal policy with the then available budget, time, and number of remaining bids. Also, the identity $V_n = \underline{U}_n$ in Proposition 5 states that the one-step dynamic programming operator applied n-many times yields the truncated version of the problem where the advertiser can bid only for the first n-many queries.

The equalities in (28) and the error bounds in (21) and (25) imply the following: given $(B, T) \in \Delta$ and a tolerance level $\varepsilon > 0$, if we fix n large so that

$$\gamma \lambda T_{\max} \cdot \min \left\{ e^{\lambda T_{\max}} \left(1 - e^{-\lambda T_{\max}} \right)^n , \ \frac{(\lambda T_{\max})^n}{(n+1)!} \right\} < \varepsilon \qquad (29)$$

and apply the policy $\left(b_i^{(n)} \right)_{i \geq 1}$, then the resulting expected total net revenue is at most ε away from the optimal expected net revenue $V(B, T)$. In other words, if n_ε is the smallest n such that (29) holds, then $(b_i^{(n_\varepsilon)})_{i \geq 1}$ is an ε-optimal bidding policy.

Corollary 2 *Because $V_n = \underline{U}_n$ for all $n \geq 0$, Propositions 3 and 4 imply that $V = \underline{U}_\infty$; namely, that V is the unique fixed point of D in $\mathcal{C}(\Delta)$.*

We already discuss the growth of V in the T variable in Remark 4. The following corollary follows from the identity $V = \underline{U}_\infty$ and the bounds in (24) in Remark 6; see also Remark 2.

Corollary 3 *When the display-and-click probability functions are all Lipschitz continuous with a Lipschitz constant α_G, we have the growth condition in (24) for the value function V.*

It should be noted that when display-and-click probability functions are not all Lipschitz continuous (everywhere), the mapping $B \mapsto V(B, T)$ may not be Lipschitz continuous (everywhere) either. To give a counterexample, let us consider a simple case with a single keyword (i.e., $K = 1$) and $G_1(b) = \sqrt{b}$ for $b \leq B_{\max} < 1$. For a bidding policy in which $b_1 = B$ and $b_i = 0$ for all $i \geq 2$, we have

$$V(B, T) \geq \mathbb{E}\left[1_{\{\tau_1 \leq T\}} (W_1 - R_1 B) \sqrt{B} \right] = \mathbb{E}\left[1_{\{\tau_1 \leq T\}} (\mu_1 - \rho_1 B) \sqrt{B} \right]$$
$$= (\mu_1 - \rho_1 B) \sqrt{B} (1 - e^{-\lambda T}),$$

which yields $\frac{V(B,T) - V(0,T)}{B} \geq \frac{\mu_1 - \rho_1 B}{\sqrt{B}} (1 - e^{-\lambda T})$, where the lower bound goes to $+\infty$ as $B \searrow 0$ for any $T > 0$.

The next proposition concludes this section with an optimal bidding policy for the main problem in (5). As in the notations for the bidding policy described in (27) and its budget process, we suppress for notational convenience the dependence of the optimal bidding policy and the corresponding budget process on the starting point (B, T) in Δ.

Proposition 6 *For $(B, T) \in \Delta$, let us extend the definition in (27) for $n = \infty$ and introduce the bidding policy $b^{(\infty)} = \left(b_i^{(\infty)} \right)_{i \geq 1}$ with its budget process $B^{(\infty)}$ (evolving according to (9)) recursively as*

$$b_i^{(\infty)} := b_{\kappa_i}^*[\underline{U}_\infty]\left(B_{\tau_{i-1}}^{(\infty)}, T - \tau_i \right) \cdot 1_{\{\tau_i \leq T\}}, \qquad for \ i \geq 1, \qquad (30)$$

again in terms of $b^*[\cdot](\cdot, \cdot)$ given in (14), and with $\tau_0 = 0$ for $i = 1$. The policy $b^{(\infty)}$ attains the supremum in (5); that is,

$$V(B, T) = \mathbb{E}\left[\sum_{i \geq 1} \left(W_i - b_i^{(\infty)} R_i\right) Z_i 1_{\{\tau_i \leq T\}}\right]. \tag{31}$$

5 Computing the value function

The successive approximations of Sect. 4, although useful in establishing the optimality of the policy in (30), is a computationally expensive way to obtain the value function. In this section, we show that the value function can be characterized as the unique solution of a differential equation with a proper boundary condition. Hence, a much faster finite difference scheme can be employed to compute the value function.

To that end, for $h \leq T \leq T_{\max}$, the expected total net revenue given on the right hand side of (31) can be decomposed as the sum of the revenue collected over the interval $[0, h]$ and that obtained over $(h, T]$. N_h denoting the number of queries over the interval $(0, h]$, the latter revenue can be written as

$$\mathbb{E}\left[\mathbb{E}\left[\sum_{i=N_h+1}^{\infty} \left(W_i - b_i^{(\infty)} R_i\right) Z_i 1_{\{\tau_i \leq T\}} \,\Big|\, \mathcal{F}_h\right]\right] = \mathbb{E}\, V\left(B_{N_h}^{(\infty)}, T - h\right),$$

where the equality is due to the Markov property. Therefore, the equality in (31) can be rewritten as

$$V(B, T) = \mathbb{E}\left[\sum_{i=1}^{N_h} \left(W_i - b_i^{(\infty)} R_i\right) Z_i + V\left(B_{N_h}^{(\infty)}, T - h\right)\right]. \tag{32}$$

For small h, using to the infinitesimal arrival probabilities of Poisson process and the distributions of W_1, R_1, and Z_1 (see also the definition of $b_1^{(\infty)}$ in (30)), we obtain

$$V(B, T) = (1 - \lambda h + o(h))V(B, T - h)$$

$$+ (\lambda h + o(h)) \sum_{k=1}^{K} \frac{\lambda_k}{\lambda}\Big\{\left(1 - G_k\left(b_k^*[V](B, T - h)\right)\right) V(B, T - h) + G_k\left(b_k^*[V](B, T - h)\right)$$

$$\times \left[\mu_k - b_k^*[V](B, T - h)\, \rho_k + \int_0^1 V\left(B - rb_k^*[V](B, T - h), T - h\right) H_k(\mathrm{d}r)\right]\Big\} + o(h)$$

$$= V(B, T - h) + h \sum_{k=1}^{K} \lambda_k\, G_k\left(b_k^*[V](B, T - h)\right) \cdot \Big[-V(B, T - h)$$

$$+ \mu_k - b_k^*[V](B, T - h)\, \rho_k + \int_0^1 V\left(B - rb_k^*[V](B, T - h), T - h\right) H_k(\mathrm{d}r)\Big] + o(h)$$

$$= V(B, T - h) + h \sum_{k=1}^{K} \lambda_k\, M_k^*[V](B, T - h) + o(h), \tag{33}$$

Step 0: For N_B, N_T large, let $\delta_B = B_{\max}/N_B$ and $\delta_T = T_{\max}/N_T$ be the grid sizes. Also let $V(j\delta_B, 0) = 0$ for $0 \le j \le N_B$, $V(0, \ell\delta_B) = 0$ for $0 \le \ell \le N_T$, and set $\ell = 1$.

Step 1: For $1 \le j \le N_B$ compute

$$V(j\delta_B, \ell\delta_T) = V(j\delta_B, (\ell-1)\delta_T) + \delta_T \sum_{k=1}^{K} \lambda_k \, M_k^*[V]\big(j\delta_B, (\ell-1)\delta_T\big),$$

where $M_k^*[V]\big(j\delta_B, (\ell-1)\delta_T\big)$ is computed via a one-dimensional numerical optimization algorithm over the interval $[0, j\delta_B \wedge (\mu_k/\rho_k)]$ for $1 \le k \le K$.

Step 2: Replace ℓ with $\ell+1$. If $\ell \le N_T$, go to Step 1; otherwise, stop.

Fig. 1 Numerical algorithm to compute the value function in (5)

where the last equality is due to the definition of $M_k^*[\cdot]$ given in (12). Recall that both V and $M_k^*[V]$ for each $1 \le k \le K$ are continuous on Δ. Therefore, after subtracting $V(B, T-h)$ from both sides in (33), dividing by h, and letting $h \to \infty$ we obtain

$$\frac{\partial V(B,T)}{\partial T} = \sum_{k=1}^{K} \lambda_k \, M_k^*[V](B,T), \tag{34}$$

which also shows that $V(\cdot, \cdot)$ is continuously differentiable in T.

Proposition 7 below verifies that the equation in (34) and the boundary condition $V(B, 0) = 0$ uniquely characterize the value function in $\mathcal{C}(\Delta)$. We can use this result and employ a finite-difference algorithm to numerically compute the value function V on Δ. That is, we discretize the space Δ, and starting with $V(\cdot, 0)$ we compute $V(\cdot, \delta_T)$, $V(\cdot, 2\delta_T)$, ..., where δ_T denotes the step length in the variable T. The details are given in Fig. 1.

Proposition 7 *If f is a function in $\mathcal{C}(\Delta)$, which is continuously differentiable in T and solves (34) with the condition $f(B, 0) = 0$ for all $B \in [0, B_{max}]$, then $f = V$.*

6 Numerical examples

In this last section, we illustrate the methods we developed so far on an example, where an advertiser wants to bid for ten keywords whose characteristics are given in Table 1. Following Cholette et al. (2012) and Dayanik and Parlar (2013), we consider a model in which, for a keyword $k \le 10$, a given bid amount b yields the following beta density

$$Q_{k,b}(d\ell) = \frac{\Gamma(a_k + b)}{\Gamma(a_k)\Gamma(b)} \ell^{a_k}(1-\ell)^{b-1} \, d\ell, \qquad \ell \in [0, 1], \tag{35}$$

for the location on the search page, where $\ell = 0$ ($\ell = 1$) corresponds to the top (bottom) of the page. A similar model is also used in Küçükaydın et al. (2020). In (35), Γ denotes the gamma function $\Gamma(\alpha) = \int_0^\infty t^{\alpha-1}e^{-t}dt$, and a_k represents the

Table 1 Keywords characteristics: arrival rates ($\vec{\lambda}$), expected revenues upon click ($\vec{\mu}$), average bids of the competitors ($\vec{a}$)

Keywords ($K = 10$)									
1	2	3	4	5	6	7	8	9	10
$\vec{\lambda}$ 50	250	100	150	200	50	80	60	90	70
$\vec{\mu}$ 50	20	30	25	20	20	25	10	20	30
$\vec{a}$ 20	100	30	10	50	35	45	40	50	60

average bid of the competitors for the k'th keyword. Conditioned on the location ℓ, the probability of click is given by $(1 - \ell)^m$ where we set $m = 0.8$ for all keywords. In this setup, the display-and-click probability function can easily be computed as

$$G_k(b) = \int_0^1 (1 - \ell)^m Q_{k,b}(d\ell) = \frac{\Gamma(a_k + b)}{\Gamma(b)} \frac{\Gamma(m + b)}{\Gamma(m + a_k + b)}, \qquad b \geq 0.$$

When a_k is an integer as in our setting, it follows from the recursion $\Gamma(\alpha) = (\alpha - 1)\Gamma(\alpha - 1)$ that

$$\frac{\Gamma(a_k + b)}{\Gamma(b)} \frac{\Gamma(m + b)}{\Gamma(m + a_k + b)} = \frac{(a_k + b - 1) \cdots b}{(m + a_k + b - 1) \cdots (m + b)} = \prod_{j=0}^{a_j - 1} \left(1 - \frac{m}{j + m + b}\right).$$

For all keywords, we use the discrete uniform distribution on $\{0.90, 0.91, \ldots, 0.99\}$ for the distributions H_k's giving the price discount effect; see (1).

Table 1 shows that keywords 2 and 5 are very popular (they have the highest search frequencies, 250 and 200 searches per day), while keyword 1 brings the highest expected sales revenue ($50). The lowest expected sales revenue ($10) is generated by keyword 8. The most frequently searched keyword; namely, keyword 2, also receives the highest average bid amount ($100) from the competitors, while the lowest average amount ($10) is bid by the competitors for keyword 4.

We calculated the optimal value and policy functions for this problem using the numerical algorithm in Fig. 1. Figure 2 shows the color-coded level plots of optimal bidding policy functions $b_k^*[V](B, T), k = 1, \ldots, 10$, of ten keywords when we start the day with a budget of $B_0 = \$2500$. Whenever a keyword search is done, we can see the optimal bid amount for that keyword from the corresponding plot at the current values of remaining time T and budget B. For example, if $T = 1/5$th of the day and $B = \$2000$ remained at the time of a keyword search was just done, then the optimal bid amounts will be about $15, $8, and $5 for keywords 1, 2, and 8, respectively. Note that the bid amounts for keyword 1 quickly rise with available budget and time. This is followed by keywords 3,4, and 10, and the slowest increase is with keyword 8. When we compare the values $(\mu_k - b_{\infty,k}\rho_k)G_k(b_{\infty,k})$, $k \leq 10$, for the keywords, we observe a similar ordering. That is, we numerically observe the intuitive result that the optimal actions become gradually myopic as the budget and time increase. The values $b_{\infty,k}$, $k \leq 10$, are reported in Table 2 below.

The perspective and level plots of value function are drawn in the upper row of Fig. 3. At every fixed remaining time T, the expected total net revenue $V(B, T)$ increases
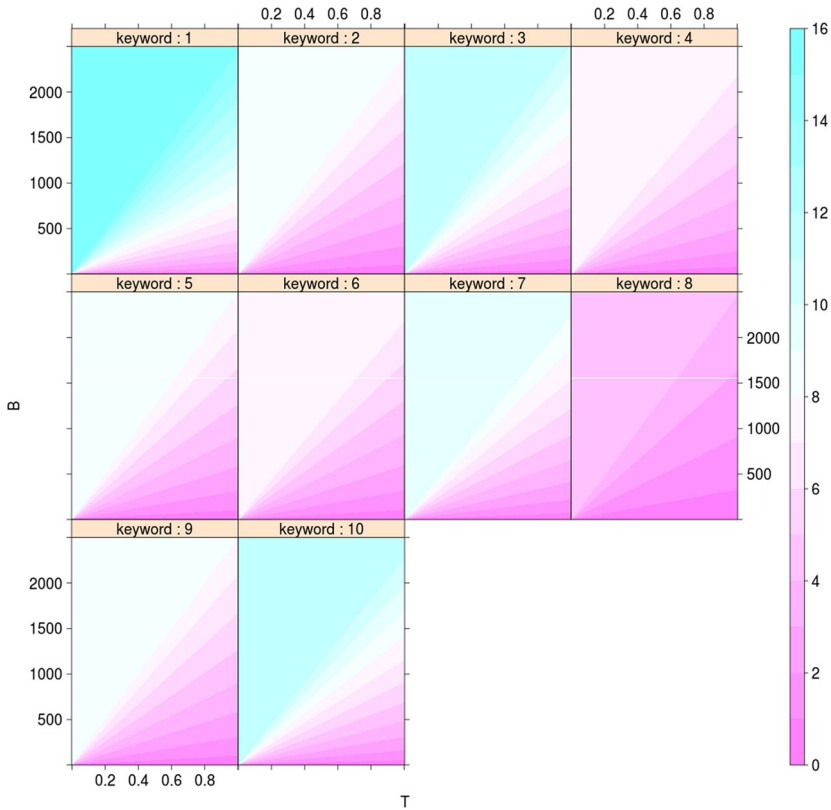
Fig. 2 Optimal bidding policy for ten keywords with the properties as in Table 1 and with a budget of $2500 at the beginning of a typical day. Horizontal axis shows the remaining fraction T of a day, and vertical axis shows the remaining advertisement budget B. Optimal bid amount for a keyword at a given remaining time and budget is color-coded according to the color key on the right and turns out to be between $0 and $16 in this example

Table 2 $b_{\infty,k}$ values for the keywords

k	1	2	3	4	5	6	7	8	9	10
$b_{\infty,k}$	16.29	9.02	11.68	8.18	8.65	8.37	10.48	4.49	8.65	12.69

with B and converges to $\gamma\lambda T = 4.47003 \times 1100 \times T = 4917.07 \times T$; see (8) in Remark 3 and the last plot in Fig. 3.

Starting with a small budget lowers the bid amounts and potentially reduces the internet traffic to product pages. A low starting budget may also cause the advertisements to completely disappear early during the day. On the other hand, starting with very large amounts may not be effective, either, because the fraction of the budget beyond what is needed to attract the total internet traffic for all of those keywords is likely to sit idle in the account. Hence, it is important to select the budget B carefully.

Fig. 3 Value function $V(B, T)$: perspective (left) and level plots (middle), and cross-sections at fixed B values (right). Last plot shows that as daily budget increases, the maximum expected net reward converges, as described in Remark 3, to the straight line $V(\infty, T) = 4917.07 \times T$

Table 3 Alternative bidding policies (left) and performances of the alternative policies. The shaded row corresponds to the daily advertisement budget $B = 2500$ beyond which additional budget is unproductive

		B	$V(B, 1)$	$\frac{\Delta V}{\Delta B}$	**Even**	$\propto \lambda\mu$	$\propto \frac{\lambda\mu}{a}$
Optimal policy		50	1390.39	962%	20.25%	15.13%	9.98%
$V(B, 1) \equiv$ 10-keywords		100	1871.10	544%	17.00%	12.60%	7.97%
$\equiv$ single budget		250	2688.26	294%	14.04%	9.95%	6.02%
1-Keyword optimal		500	3423.76	149%	12.27%	8.26%	5.33%
Even $\equiv$ budget split		1000	4171.24	72%	10.37%	6.65%	5.16%
$\equiv$ evenly		1500	4533.81	34%	8.80%	5.53%	5.20%
$\lambda\mu \equiv$ budget split		2000	4706.51	11%	7.08%	4.35%	4.63%
$\equiv \propto \lambda\mu$		2500	4761.91	0%	5.03%	2.83%	3.27%
$\lambda\mu/a \equiv$ budget split		3000	4761.97	0%	3.04%	1.35%	1.79%
$\equiv \propto \lambda\mu/a$		4000	4761.97	0%	1.01%	0.20%	0.49%
		5000	4761.97	NA	0.26%	0.02%	0.12%

The first two columns of Table 3 show the maximum expected total net revenues $V(B, 1)$ corresponding to starting budget values $B = 50, 100, \dots, 5000$ using the same ten keywords of Table 1. Note that $V(B, 1)$ increases in significant amounts with increasing B initially and flattens for $B \geq 2500$. Those figures suggest that an advertisement budget should be chosen somewhere between 2000 and 2500. We can devise a more precise guideline based on the third column of the same Table 3, which reports

$$\frac{\Delta V}{\Delta B} = \frac{V(B + \Delta B, 1) - V(B, 1)}{\Delta B} \times 100\%.$$

For example, since $B = 2000$, $V(2000, 1) = 4706.51$, $V(2500, 1) = 4761.91$, and $\Delta B = 2500 - 2000 = 500$, we find $\Delta V/\Delta B = (4761.91 - 4706.51)/500 = 11\%$.
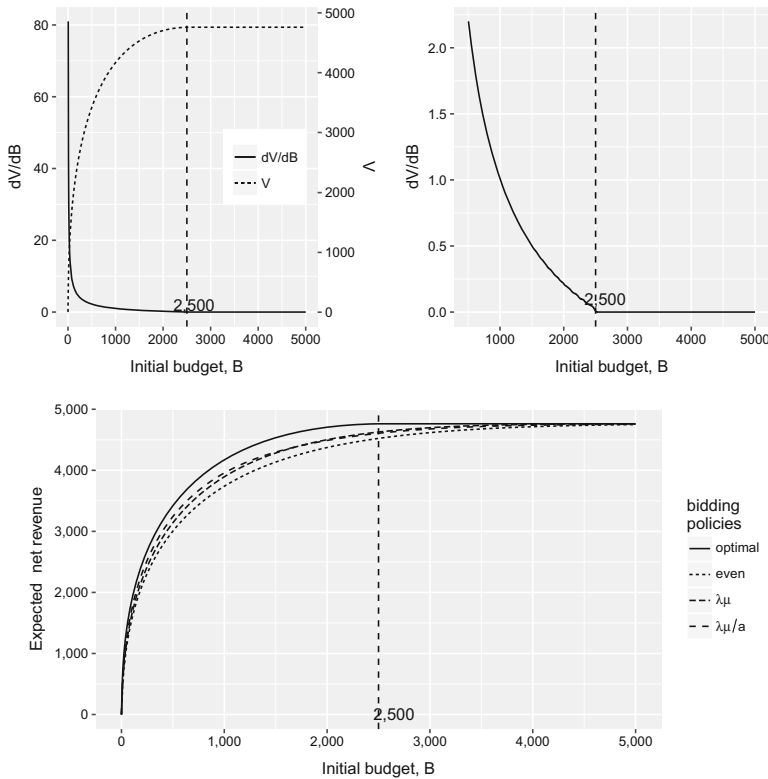
Fig. 4 Selection of the initial budget (top) and comparison of bidding policies (bottom). On the upper left plot, the expected total net revenue (on the vertical axis on the right) increases, while the rate of increase (on the vertical axis on the left) decreases very fast. We zoomed the picture on the upper right plot to show better that the rate vanishes near $B = 2500$, which we set as the daily advertisement budget. The bottom plot gives the expected daily net revenues of optimal and heuristic bidding policies of Table 3 (left). When we start with the initial budget $B = 2500$, to close up the revenue gap between optimal and heuristic bidding policies, one needs to nearly double the daily budget to 5000; see also Table 3 (right)

This ratio shows roughly the daily marginal rate of return on the additional $\Delta B = 500$ investment into search-based advertisement budget B. The top row of Fig. 4 gives the full story: it gives the ratio $\Delta V / \Delta B$ at every budget B value with $\Delta B = 1$. That is, we consider adding one additional unit budget to see its effect. In practice, a good rule of thumb would be to keep incrementing the budget until some target internal rate of return is reached. In Table 3 and Fig. 4, we observe that any additional budget beyond $B = 2500$ is simply unproductive. Hence, one should not go beyond that level. If there is no alternative investment opportunity (to make a rate of return comparison), $B = 2500$ would be a good choice as a daily search-based advertisement budget to bid for those ten keywords in Table 1.

Along with the optimal multi-keyword bidding policy, we also calculate expected total net revenues of three one-keyword heuristic bidding policies described in Table 3 (left). In all of those heuristics, the budget is split at the very beginning and then used exclusively for each of ten keywords (with their one-keyword optimal policies)

Fig. 5 Counter example for the monotonicity of the optimal bid amounts in B: plots of the display-and-click probability function $b \mapsto G(b)$ (left), the mapping $B \mapsto V(B, 1)$ (middle), optimal bid amounts $B \mapsto b^*[V](B, 1)$ (right)

throughout the day. In the "even" heuristic, the budget is evenly distributed among ten keywords. In the "$\lambda\mu$" and "$\lambda\mu/a$" heuristics, the budget is split among the keywords proportional to $\lambda\mu$ and $\lambda\mu/a$, respectively. The former heuristic tries to capture each keyword's capability to generate revenues in the absence of any competitors, while the latter heuristic tries to also take into account the strength of competition between rivals. Table 3 (right) provides the percentage shortages

$$\frac{V(B, 1) - \text{expected revenue of the heuristic}}{V(B, 1)} \times 100\%$$

for each of heuristics "even", "$\propto \lambda\mu$", and "$\propto \lambda\mu/a$".

We observe that the multi-keyword optimal policy differs significantly from the others when the daily advertisement budget is tight, and the differences diminish as the budget increases as expected. Three heuristic bidding policies are ordered as $\propto \lambda\mu/a, \propto \lambda\mu$, "even" from the best to the worse, respectively, for most of the budget values in Table 3. The bottom row of Fig. 4 gives a more complete comparison of optimal and heuristic bidding policies.

Recall that additional budget beyond \$2500 did not increase the expected net revenue much under optimal bidding policy. With this budget at hand, the optimal strategy earns more than any heuristic by a margin between 2.83% and 5.01% every day; see the shaded row of Table 3 (right). We see the corresponding gaps between solid and dashed curves along the vertical dashed line at $B = 2500$ in Fig. 4 (bottom). If one is prepared to invest a higher advertisement budget into any of those heuristic methods to close up the performance gaps, the last row of Table 3 (right) tells us that, to reduce the this gap to less than 1%, one needs a budget around \$5000, which is *twice* the optimal budget amount. Likewise, Fig. 4 (bottom) shows that expected net revenues from heuristics increase slowly after $B = 2500$, and in order to close up the gap, the budget should nearly be doubled to 5000.

If, on the other hand, the advertisement budget has to be cut back, for example, for the lack of sufficient funds, then the margin between optimal and heuristic bidding policies can be as large as 10–20%, and those are the significant gains of the optimal bidding strategy over those three heuristic bidding policies.

Lastly, we conclude with an example in which the bid amounts are not always monotone in the remaining budget. For that, we consider a problem with a single keyword ($K = 1$) for which $\lambda = 50$ and $\mu = 10$. We assume the same H_k function as in the previous example with ten keywords. For the display-and-click probability function, we take a logistic function of the form

$$G(b) = \frac{0.5}{1 + 100b^{-5}}, \qquad b \geq 0,$$

shown in the first plot in Fig. 5. In this example, we observe a non-monotone behavior in the optimal bid amount $B \mapsto b^*[V](B, 1)$; see the last plot in the figure. The plot of the function $b \mapsto G(b)$ shows that small bid amounts bring higher marginal increments in the display-and-click-probabilities. Considering that the advertiser has only a small number of bidding opportunities when the budget is relatively tight, at certain levels, it becomes preferable to decrease the bid amount and leave some amount for later in order to potentially benefit more from this region of high marginal increments. We observe a wavy behavior in the slope of the value function again especially over the region where the budget is tight; see the second plot in the same figure for $B \mapsto V(B, 1)$.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s00186-022-00803-y.

Appendix: Supplementary proofs

Proof of the equality in (7) For a given admissible bidding policy, since $\sum_{n\geq 1} b_n R_n Z_n 1_{\{\tau_n \leq T\}}$ is bounded by B, we have

$$\mathbb{E}\left[\sum_{n\geq 1}(W_n - b_n R_n)Z_n 1_{\{\tau_n \leq T\}}\right] = \mathbb{E}\left[\sum_{n\geq 1} W_n Z_n 1_{\{\tau_n \leq T\}}\right] - \mathbb{E}\left[\sum_{n\geq 1} b_n R_n Z_n 1_{\{\tau_n \leq T\}}\right]$$

$$= \sum_{n\geq 1}\mathbb{E}[W_n Z_n 1_{\{\tau_n \leq T\}}] - \sum_{n\geq 1}\mathbb{E}[b_n R_n Z_n 1_{\{\tau_n \leq T\}}],$$

$$(36)$$

where the second line follows by the monotone convergence theorem applied to each expected random sum separately. Note that

$$\mathbb{E}[W_n Z_n 1_{\{\tau_n \leq T\}}] = \mathbb{E}[1_{\{\tau_n \leq T\}}\mathbb{E}[W_n Z_n \mid \mathcal{F}_{\tau_n}]] = \mathbb{E}[1_{\{\tau_n \leq T\}}\mu_{\kappa_n} G_{\kappa_n}(b_n)], \quad (37)$$

due to the conditional independence of W_n, the conditional expectations in (3), and the conditional distributions in (1). Similarly, we write

$$\mathbb{E}[b_n R_n Z_n 1_{\{\tau_n \leq T\}}] = \mathbb{E}[1_{\{\tau_n \leq T\}}b_n\mathbb{E}[R_n Z_n \mid \mathcal{F}_{\tau_n}]] = \mathbb{E}[1_{\{\tau_n \leq T\}}b_n \rho_{\kappa_n} G_{\kappa_n}(b_n)], \quad (38)$$

thanks to (1) and (2). Using now (37-38) in (36) yields

$$
\begin{aligned}
\mathbb{E}&\Big[\sum_{n \geq 1} (W_n - b_n R_n) Z_n 1_{\{\tau_n \leq T\}} \Big] \\
&= \sum_{n \geq 1} \mathbb{E}[1_{\{\tau_n \leq T\}} \mu_{\kappa_n} G_{\kappa_n}(b_n)] - \sum_{n \geq 1} \mathbb{E}[1_{\{\tau_n \leq T\}} b_n \rho_{\kappa_n} G_{\kappa_n}(b_n)] \\
&= \mathbb{E}\Big[\sum_{n \geq 1} 1_{\{\tau_n \leq T\}} \mu_{\kappa_n} G_{\kappa_n}(b_n) \Big] - \mathbb{E}\Big[\sum_{n \geq 1} 1_{\{\tau_n \leq T\}} b_n \rho_{\kappa_n} G_{\kappa_n}(b_n) \Big] \\
&= \mathbb{E}\Big[\sum_{n \geq 1} 1_{\{\tau_n \leq T\}} (\mu_{\kappa_n} - b_n \rho_{\kappa_n}) G_{\kappa_n}(b_n) \Big]
\end{aligned}
\tag{39}
$$

establishing (7). In (39), the second equality is by the monotone convergence theorem (applied to each expectation), and the last equality follows simply by the boundedness of $\sum_{n \geq 1} 1_{\{\tau_n \leq T\}} b_n \rho_{\kappa_n} G_{\kappa_n}(b_n)$ (by B). $\qquad \square$

Proof of Proposition 1 Monotonicity of $\mathcal{D}[f]$ in f is obvious, and the non-negativity of $D[f_1]$ follows after taking $b_1 = 0$ in (10). To prove the upper bound $\mathcal{D}[f_2](\cdot, T) \leq \gamma \lambda T$, we observe that

$$
\begin{aligned}
\mathcal{D}[f](B,T) &\leq \sup_{b_1 \in \mathcal{F}_{\tau_1}} \mathbb{E}\, 1_{\{\tau_1 \leq T\}} \Big[(W_1 - b_1 R_1) Z_1 + \gamma \lambda (T - \tau_1) \Big] \\
&= \sup_{b_1 \in \mathcal{F}_{\tau_1}} \mathbb{E}\, 1_{\{\tau_1 \leq T\}} \Big[(\mu_{\kappa_1} - b_1 \rho_{\kappa_1}) G_{\kappa_1}(b_1) + \gamma \lambda (T - \tau_1) \Big] \\
&\leq \mathbb{E}\, 1_{\{\tau_1 \leq T\}} \Big[\gamma + \gamma \lambda (T - \tau_1) \Big] = \gamma \int_0^T \lambda e^{-\lambda t} [1 + \lambda (T - t)] dt = \gamma \lambda T,
\end{aligned}
$$

in which the second line follows by conditioning on $\mathcal{F}_{\tau_1}$, the inequality in the third line is by related arguments on the mapping $b \mapsto (\mu_k - b \rho_k) G_k(b)$ and the definition of γ in Remark 3, and the very last equality is simply by integration (by parts). $\qquad \square$

Proof of Proposition 2 For notational convenience, let us define for $1 \leq k \leq K$

$$
\begin{aligned}
L_k[f](b, B, T) &:= f(B, T) + M_k[f](b, B, T) \quad \text{and} \\
L_k^*[f](B, T) &:= f(B, T) + M_k^*[f](B, T),
\end{aligned}
$$

with M_k and M_k^* defined in (12-13). It is easy to verify that the mappings $B \mapsto L_k^*[f](B, \cdot)$ and $T \mapsto L_k^*[f](\cdot, T)$ are non-decreasing under the given assumptions (on the monotonicity of f in its arguments). Since

$$
\mathcal{D}[f](B, T) = \int_0^T \sum_{k=1}^K \lambda_k e^{-\lambda u} L_k^*[f](B, T - u) du,
$$

it follows that $B \mapsto \mathcal{D}[f](B, \cdot)$ and $T \mapsto \mathcal{D}[f](\cdot, T)$ are again non-decreasing. This proves the non-negativity of the differences in (16) and (17).

The prove the upper bound in (16), we note that, for $T_1 < T_2$, we have

$$\mathcal{D}[f](B, T_2) - \mathcal{D}[f](B, T_1) = (e^{-\lambda T_2} - e^{-\lambda T_1}) \int_0^{T_1} \sum_{k=1}^K \lambda_k e^{\lambda u} L_k^*[f](B, u) du$$

$$+ e^{-\lambda T_2} \int_{T_1}^{T_2} \sum_{k=1}^K \lambda_k e^{\lambda u} L_k^*[f](B, u) du$$

$$\leq \int_{T_1}^{T_2} \sum_{k=1}^K \lambda_k L_k^*[f](B, u) du. \tag{40}$$

It is easy to verify that $L_k^*[f](B, T) \leq \mu_k + \|f\| \leq \bar{\mu} + \|f\|$. Hence, we obtain $\mathcal{D}[f](B, T_2) - \mathcal{D}[f](B, T_1) \leq \lambda(\bar{\mu} + \|f\|)(T_2 - T_1)$.

To establish the second claim, let $B_1 < B_2$ be two budget levels, and for fixed T, let $b_2 = b_k^*[f](B_2, T)$ denote the maximum bid in (14) for a fixed k with the budget level B_2. Note that $b_1 = \frac{B_1}{B_2} b_2 < b_2$ is a feasible bid for the budget B_1, and we have $b_2 - b_1 = \frac{(B_2 - B_1)}{B_2} b_2 \leq B_2 - B_1$. Then, we write

$$L_k^*[f](B_2, T) = (1 - G_k(b_2)) f(B_2, T) + G_k(b_2) \Big[\mu_k - \rho_k b_2$$

$$+ \int_0^1 f(B_2 - r b_2, T) H_k(dr) \Big]$$

$$\leq (1 - G_k(b_2))(f(B_1, T) + \alpha_f (B_2 - B_1))$$

$$+ G_k(b_2) \Big[\mu_k - \rho_k b_1 + \int_0^1 f(B_1 - r b_1, T) H_k(dr) + \alpha_f (B_2 - B_1) \Big]$$

$$= \alpha_f (B_2 - B_1) + (1 - G_k(b_2)) f(B_1, T)$$

$$+ G_k(b_2) \Big[\mu_k - \rho_k b_1 + \int_0^1 f(B_1 - r b_1, T) H_k(dr) \Big]$$

$$\leq \alpha_f (B_2 - B_1) + L_k[f](b_1, B_1, T) + \alpha_G(b_2 - b_1)(\bar{\mu} + \|f\|)$$

$$\leq \alpha_f (B_2 - B_1) + L_k^*[f](B_1, T) + \alpha_G(B_2 - B_1)(\bar{\mu} + \|f\|)$$

$$= L_k^*[f](B_1, T) + (B_2 - B_1)\big[\alpha_f + \alpha_G(\bar{\mu} + \|f\|)\big]. \tag{41}$$

Next, we consider the difference

$$\mathcal{D}[f](B_2, T) - \mathcal{D}[f](B_1, T) = e^{-\lambda T} \int_0^T \sum_{k=1}^K \lambda_k e^{\lambda u} \big[L_k^*[f](B_2, u)$$

$$- L_k^*[f](B_1, u) \big] du,$$

and using (41), we obtain $\mathcal{D}[f](B_2, T) - \mathcal{D}[f](B_1, T) \leq$

$$e^{-\lambda T} \int_0^T \sum_{k=1}^K \lambda_k e^{\lambda u} (B_2 - B_1)\big[\alpha_f + \alpha_G(\bar{\mu} + \|f\|)\big] du$$

$$\leq (B_2 - B_1)(1 - e^{-\lambda T_{\max}})\big[\alpha_f + \alpha_G(\bar{\mu} + \|f\|)\big]$$

giving us the upper bound in (17). □

Proof of Proposition 3 Because $\underline{U}_0(\cdot, \cdot) = 0 \leq \overline{U}_0(B, T) = \gamma\lambda T$, Proposition 1 implies that $0 \leq \underline{U}_1 \leq \overline{U}_1 \leq \gamma\lambda T$. Let us now assume that $\underline{U}_{n-1} \leq \underline{U}_n \leq \overline{U}_n \leq \overline{U}_{n-1}$ for some $n \geq 1$. Then, again by Proposition 1, we have $\mathcal{D}[\underline{U}_{n-1}] \leq \mathcal{D}[\underline{U}_n] \leq \mathcal{D}[\overline{U}_n] \leq \mathcal{D}[\overline{U}_{n-1}]$, and together with the induction hypothesis, this yields $\underline{U}_{n-1} \leq \underline{U}_n \leq \underline{U}_{n+1} \leq \overline{U}_{n+1} \leq \overline{U}_n \leq \overline{U}_{n-1}$. Hence, by induction we conclude that $(\underline{U}_n(\cdot, \cdot))_{n\geq 0}$ and $(\overline{U}_n(\cdot, \cdot))_{n\geq 0}$ are non-decreasing and non-increasing, respectively, and the collection $(\underline{U}_n(\cdot, \cdot))_{n\geq 0}$ is bounded from above by $(\overline{U}_n(\cdot, \cdot))_{n\geq 0}$.

Let $\underline{U}_\infty$ denote the pointwise limit of the monotone sequence $(\underline{U}_n)_{n\geq 0}$. Because the operator $\mathcal{D}$ is a contraction mapping, $\|\underline{U}_{\ell+1} - \underline{U}_\ell\| = \|\mathcal{D}[\underline{U}_\ell] - \mathcal{D}[\underline{U}_{\ell-1}]\| \leq (1 - e^{-\lambda T_{\max}})\|\underline{U}_\ell - \underline{U}_{\ell-1}\| \leq \ldots \leq (1 - e^{-\lambda T_{\max}})^\ell \|\underline{U}_1\|$, using which we obtain

$$\|\underline{U}_{m+n} - \underline{U}_n\| \leq \sum_{\ell=n}^{m+n-1} \|\underline{U}_{\ell+1} - \underline{U}_\ell\| \leq \|\underline{U}_1\| \sum_{\ell=n}^{m+n-1} (1 - e^{-\lambda T_{\max}})^\ell \text{ for every } m, n \geq 0,$$

and this gives $0 \leq \underline{U}_\infty(B, T) - \underline{U}_n(B, T) = \lim_{m\to\infty} \underline{U}_{m+n}(B, T) - \underline{U}_n(B, T) \leq \|\underline{U}_1\| \lim_{m\to\infty} \sum_{\ell=n}^{m+n-1}(1 - e^{-\lambda T_{\max}})^\ell \leq \gamma\lambda T_{\max} \sum_{\ell=n}^{\infty}(1 - e^{-\lambda T_{\max}})^\ell = \gamma\lambda T_{\max}(1 - e^{-\lambda T_{\max}})^n e^{\lambda T_{\max}}$. Since this is true for all $(B, T) \in \Delta$, we have $\|\underline{U}_\infty - \underline{U}_n\| \leq \gamma\lambda T_{\max}(1 - e^{-\lambda T_{\max}})^n e^{\lambda T_{\max}}$, which proves that the convergence of $\underline{U}_n$ to $\underline{U}_\infty$ is uniform on Δ as $n \to \infty$. In the arguments above, if we replace $\underline{U}$ with $\overline{U}$, then we obtain the same upper bound for $\overline{U}_n(B, T) - \overline{U}_\infty(B, T) \geq 0$. This shows that the $\overline{U}_n$'s converge to $\overline{U}_\infty(B, T)$ also uniformly with the same error bound. Because $\overline{U}_n$'s and $\underline{U}_n$'s are continuous, their uniform limits $\overline{U}_\infty$ and $\underline{U}_\infty$, respectively, are also continuous on Δ.

The uniform convergence of $(\underline{U}_n)_{n\geq 0}$ and $(\overline{U}_n)_{n\geq 0}$ also imply that as $n \to \infty$, we have $M_k^*[\overline{U}_n](\cdot, \cdot) \to M_k^*[\overline{U}_\infty](\cdot, \cdot)$ and $M_k^*[\underline{U}_n](\cdot, \cdot) \to M_k^*[\underline{U}_\infty](\cdot, \cdot)$ for each $1 \leq k \leq K$, and they are bounded from above by $\max_k M_k^*[\overline{U}_0](B_{\max}, T_{\max}) < \infty$. Then, by bounded convergence theorem we obtain

$$\underline{U}_\infty(B, T) = \lim_n \underline{U}_{n+1}(B, T) = \lim_n \mathcal{D}[\underline{U}_n](B, T) = \lim_n \mathbb{E}\, 1_{\{\tau_1 \leq T\}}\Big[\underline{U}_n(B, T - \tau_1)$$
$$+ M_{\kappa_1}^*[\underline{U}_n](B, T - \tau_1)\Big]$$
$$= \mathbb{E}\, 1_{\{\tau_1 \leq T\}}\Big[\underline{U}_\infty(B, T - \tau_1) + M_{\kappa_1}^*[\underline{U}_\infty](B, T - \tau_1)\Big] = \mathcal{D}[\underline{U}_\infty](B, T),$$

which shows that $\underline{U}_\infty$ is a fixed point of $\mathcal{D}$. Replicating the arguments above with $\overline{U}_\infty = \lim_n \overline{U}_n$, we observe that $\overline{U}_\infty$ is also a fixed point of the operator $\mathcal{D}$. Finally, the uniqueness of the fixed point (see Lemma 1) implies that $\underline{U}_\infty = \overline{U}_\infty$. □

Proof of Lemma 2 The inequality in (26) becomes an equality for $j = 1$. Assume that it holds for some $1 \leq j \leq n$, and let us prove it for $j + 1$. Note that the right hand

side of (26) can be decomposed as

$$
\mathbb{E} \sum_{i=1}^{n-j} (W_i - b_i R_i) Z_i 1_{\{\tau_i \leq T\}}
$$

$$
+ \mathbb{E} \left[1_{\{\tau_{n-j+1} \leq T\}} \Big((W_{n-j+1} - b_{n-j+1} R_{n-j+1}) Z_{n-j+1} + \underline{U}_{j-1}(B_{\tau_{n+j+1}}, T - \tau_{n-j+1}) \Big) \right].
$$

$$
\tag{42}
$$

Conditioning on $\mathcal{F}_{\tau_{n-j+1}}$ and using the conditional distributions of W_{n-j+1}, R_{n-j+1}, and Z_{n-j+1}, the second expectation in (42) above becomes

$$
\mathbb{E}\left[1_{\{\tau_{n-j+1} \leq T\}} \Big((\mu_{\kappa_{n-j+1}} - b_{n-j+1} \cdot \rho_{n-j+1}) G_{\kappa_{n-j+1}}(b_{n-j+1}) \right.
$$

$$
+ \underline{U}_{j-1}(B_{\tau_{n-j}}, T - \tau_{n-j+1})(1 - G_{\kappa_{n-j+1}}(b_{n-j+1}))
$$

$$
\left. + G_{\kappa_{n-j+1}}(b_{n-j+1}) \int_0^1 \underline{U}_{j-1}(B_{\tau_{n-j}} - r b_{n-j+1}, T - \tau_{n-j+1}) h_{\kappa_{n-j+1}}(r) dr \Big) \right]
$$

$$
= \mathbb{E}\left[1_{\{\tau_{n-j+1} \leq T\}} \Big(\underline{U}_{j-1}(B_{\tau_{n-j}}, T - \tau_{n-j+1}) \right.
$$

$$
\left. + M_{\kappa_{n-j+1}}[\underline{U}_{j-1}](b_{n-j+1}, B_{\tau_{n-j}}, T - \tau_{n-j+1}) \Big) \right]
$$

$$
\leq \mathbb{E}\left[1_{\{\tau_{n-j+1} \leq T\}} \Big(\underline{U}_{j-1}(B_{\tau_{n-j}}, T - \tau_{n-j+1}) + M^*_{\kappa_{n-j+1}}[\underline{U}_{j-1}](B_{\tau_{n-j}}, T - \tau_{n-j+1}) \Big) \right]
$$

$$
= \mathbb{E}\left[1_{\{\tau_{n-j} \leq T\}} \mathbb{E}\left[1_{\{\tau_{n-j+1} \leq T\}} \Big(\underline{U}_{j-1}(B_{\tau_{n-j}}, T - \tau_{n-j+1}) \right.\right.
$$

$$
\left.\left. + M^*_{\kappa_{n-j+1}}[\underline{U}_{j-1}](B_{\tau_{n-j}}, T - \tau_{n-j+1}) \Big) \Big| \mathcal{F}_{\tau_{n-j}} \vee \sigma(B_{\tau_{n-j}}) \right]\right]
$$

$$
= \mathbb{E}\left[1_{\{\tau_{n-j} \leq T\}} D[\underline{U}_{j-1}](B_{\tau_{n-j}}, T - \tau_{n-j}) \right]
$$

$$
\tag{43}
$$

where the last line is due to strong Markov property. Because $\mathcal{D}[\underline{U}_{j-1}] = \underline{U}_j$, combining (42-43) with the induction hypothesis yields

$$
\mathbb{E}\left[\sum_{i=1}^{n} (W_i - b_i R_i) Z_i 1_{\{\tau_i \leq T\}} \right] \leq \mathbb{E}\left[\sum_{i=1}^{n-j+1} (W_i - b_i R_i) Z_i 1_{\{\tau_i \leq T\}} \right.
$$

$$
\left. + 1_{\{\tau_{n-j+1} \leq T\}} \underline{U}_{j-1}(B_{\tau_{n+j+1}}, T - \tau_{n-j+1}) \right]
$$

$$
\leq \mathbb{E}\left[\sum_{i=1}^{n-j} (W_i - b_i R_i) Z_i 1_{\{\tau_i \leq T\}} + 1_{\{\tau_{n-j} \leq T\}} \right.
$$

$$
\left. \underline{U}_j(B_{\tau_{n-j}}, T - \tau_{n-j}) \right],
$$

$$
\tag{44}
$$

and this proves the inequality (26) for $j + 1$. Hence, by induction it holds for all $1 \leq j \leq n + 1$. $\qquad \square$

Proof of Proposition 5 For $n = 0$, we have $V_0 = \underline{U}_0 = 0$ by construction, and the summation in (28) equals zero. Hence, the claim of the proposition is obvious. Therefore, we state the proposition and prove the equalities in (28) for $n \geq 1$.

By Lemma 2, we have for every admissible policy $(b_i)_{i \geq 1}$

$$\mathbb{E}\Big[\sum_{i=1}^{n}(W_i - b_i R_i)Z_i 1_{\{\tau_i \leq T\}}\Big]$$

$$\leq \mathbb{E}\Big[\sum_{i=1}^{n-j+1}(W_i - b_i R_i)Z_i 1_{\{\tau_i \leq T\}} + 1_{\{\tau_{n-j+1} \leq T\}}\underline{U}_{j-1}(B_{\tau_{n-j+1}}, T - \tau_{n-j+1})\Big]$$

for $1 \leq j \leq n + 1$. Evaluating this inequality with $k = n + 1$ gives $\mathbb{E}\Big[\sum_{i=1}^{n}(W_i - b_i R_i)Z_i 1_{\{\tau_i \leq T\}}\Big] \leq \underline{U}_n(B, T)$ and this implies $V_n(B, T) \leq \underline{U}_n(B, T)$ for all $n \geq 1$.

We next establish the second equality in (28) for all $n \geq 1$. Because $\underline{U}_0 \equiv 0$, we have

$$\begin{aligned}
\underline{U}_1(B, T) = \mathcal{D}[\underline{U}_0](B, T) &= \sup_{b_1 \in \mathcal{F}_{\tau_1}} \mathbb{E}\Big[1_{\{\tau_1 \leq T\}}(W_1 - b_1 R_1)Z_1\Big] \\
&= \sup_{b_1 \in \mathcal{F}_{\tau_1}} \mathbb{E}\, 1_{\{\tau_1 \leq T\}} G_{\kappa_1}(b_1)\big(\mu_{\kappa_1} - b_1 \rho_{\kappa_1}\big) \\
&= \mathbb{E}\Big[1_{\{\tau_1 \leq T\}}\big(W_1 - b_{\kappa_1}^*[\underline{U}_0](B, T - \tau_1) \cdot R_1\big)Z_1\Big] \\
&\equiv \mathbb{E}\Big[1_{\{\tau_1 \leq T\}}\big(W_1 - b_1^{(1)} R_1\big)Z_1\Big],
\end{aligned}$$

and this gives the second equality in (28) for $n = 1$. Assume now that the second equality in (28) holds for some $n \geq 1$. Then we have

$$\begin{aligned}
\underline{U}_{n+1}(B, T) &= \mathcal{D}[\underline{U}_n](B, T) \\
&= \sup_{b_1 \in \mathcal{F}_{\tau_1}} \mathbb{E}\, 1_{\{\tau_1 \leq T\}}\Big[(W_1 - b_1 R_1)Z_1 + \underline{U}_n(B_{\tau_1}, T - \tau_1)\Big] \\
&= \mathbb{E}\, 1_{\{\tau_1 \leq T\}}\Big[(W_1 - \underbrace{b_{\kappa_1}^*[\underline{U}_n](B, T - \tau_1)}_{b_1^{(n+1)}} \cdot R_1)Z_1 + \underline{U}_n(B_{\tau_1}, T - \tau_1)\Big].
\end{aligned}$$

$$(45)$$

Using the induction hypothesis and the strong Markov property, we obtain

$$
\mathbb{E}\big(1_{\{\tau_1 \le T\}}\, \underline{U}_n(B_{\tau_1}, T - \tau_1)\big)
$$

$$
= \mathbb{E}\bigg(1_{\{\tau_1 \le T\}}\mathbb{E}\Big[\sum_{i=1}^{n}\big(W_{i+1} - b^*_{\kappa_{i+1}}[\underline{U}_{n-i}](B^{(n+1)}_{T_i}, T - T_{i+1})
$$

$$
\cdot R_{i+1}\big)\, Z_{i+1}1_{\{T_{i+1} \le T\}} \mid \mathcal{F}_{\tau_1} \vee \sigma(B_{\tau_1})\Big]\bigg)
\tag{46}
$$

$$
= \mathbb{E}\bigg[\sum_{i=2}^{n+1}\Big(W_i - b^*_{\kappa_i}[\underline{U}_{n+1-i}](B^{(n+1)}_{\tau_{i-1}}, T - \tau_i) \cdot R_i\Big)\, Z_i\, 1_{\{\tau_i \le T\}}\bigg]
$$

$$
\equiv \mathbb{E}\bigg[\sum_{i=2}^{n+1}\Big(W_i - b^{(n+1)}_i R_i\Big)\, Z_i\, 1_{\{\tau_i \le T\}}\bigg].
$$

Substituting (46) into (45) gives

$$
\underline{U}_{n+1}(B, T) = \mathbb{E}\bigg[\sum_{i=1}^{n+1}\Big(W_i - b^{(n+1)}_i R_i\Big)\, Z_i\, 1_{\{\tau_i \le T\}}\bigg].
$$

This proves the second equality in (28) for $n + 1$. Hence, it holds for all $n \ge 1$ by induction.

Clearly, V_n is an upper bound for the expected net revenue of any n-bid policy. Hence, combining all the arguments above, we now have, for any $n \ge 1$,

$$
V_n(B, T) \le \underline{U}_n(B, T) = \mathbb{E}\bigg[\sum_{i=1}^{n}\Big(W_i - b^{(n+1)}_i R_i\Big)\, Z_i\, 1_{\{\tau_i \le T\}}\bigg] \le V_n(B, T),
$$

and this establishes the equalities in (28). □

Proof of Proposition 6 To prove the claim, it is sufficient to establish the identity

$$
V(B, T) = \mathbb{E}\bigg[\sum_{i=1}^{n}\Big(W_i - b^{(\infty)}_i R_i\Big)\, Z_i 1_{\{\tau_i \le T\}} + 1_{\{\tau_n \le T\}}\, V\big(B^{(\infty)}_{\tau_n}, T - \tau_n\big)\bigg], \tag{47}
$$

inductively for all $n \ge 1$. When we let $n \to \infty$, the expectation of the second term converges to zero because V is bounded (see Remark 3) and $\mathbb{P}(\tau_n \le T) \to 0$. Also both $\mathbb{E}\big[\sum_{i=1}^{n} W_i \cdot R_i Z_i 1_{\{\tau_i \le T\}}\big]$ and $\mathbb{E}\big[\sum_{i=1}^{n} b^{(\infty)}_i R_i Z_i 1_{\{\tau_i \le T\}}\big]$ convergence by monotone convergence theorem to the expectations of the corresponding infinite sums, and $\sum_{i=1}^{\infty} b^{(\infty)}_i R_i Z_i 1_{\{\tau_i \le T\}} \le B$. Hence, as $n \to \infty$, the expectation of the summation in (47) converges to the right hand side in (31).

For $n = 1$, the expectation in (47) becomes $\mathcal{D}[V](B, T)$, and the equality holds since V is a fixed point of the operator $\mathcal{D}$. Assume now that the equality holds for some $n \ge 1$. Then

$$\mathbb{E}\left[\sum_{i=1}^{n+1}\left(W_i - b_i^{(\infty)}R_i\right)Z_i 1_{\{\tau_i \le T\}} + 1_{\{T_{n+1}\le T\}} V\left(B_{T_{n+1}}^{(\infty)}, T - T_{n+1}\right)\right]$$

$$= \mathbb{E}\left[\left(W_1 - b_1^{(\infty)}R_1\right)Z_1 1_{\{\tau_1 \le T\}}\right.$$

$$+ 1_{\{\tau_1 \le T\}}\mathbb{E}\left[\sum_{i=2}^{n+1}\left(W_i - b_i^{(\infty)}R_i\right)Z_i 1_{\{\tau_i \le T\}}\right.$$

$$\left.\left. + 1_{\{T_{n+1}\le T\}}V\left(B_{T_{n+1}}^{(\infty)}, T - T_{n+1}\right) \,\Big|\, \mathcal{F}_{\tau_1} \vee \sigma(B_{T_1}^{(\infty)})\right]\right]. \tag{48}$$

On the event $\{\tau_1 \le T\}$, the conditional expectation above is equal to $V\left(B_{\tau_1}^{(\infty)}, T - \tau_1\right)$ by the strong Markov property and the induction hypothesis. Hence, the right hand side of the equality in (48) becomes

$$\mathbb{E}1_{\{\tau_1 \le T\}}\left[\left(W_1 - b_1^{(\infty)}R_1\right)Z_1 + V\left(B_{\tau_1}^{(\infty)}, T - \tau_1\right)\right] = \mathcal{D}[V](B, T) = V(B, T),$$

which proves (47) for $n + 1$. Hence, (47) holds for all $n \ge 1$ by induction, and this completes the proof. □

Proof of Proposition 7 The identity $f(B, 0) = V(B, 0)$ for $T = 0$ is obvious. Therefore, we only give the proof for $T > 0$. Note that for a given admissible policy $(b_i)_{i \ge 1}$ and the corresponding budget process $\{B_t\}_{t \in [0,T]}$, the chain rule gives

$$\mathbb{E}\underbrace{f(B_T, 0)}_{0} - f(B, T)$$

$$= \mathbb{E}\left[-\int_0^T f_T(B_{t-}, T - t)dt + \sum_{i \ge 1}\left[f(B_{\tau_i}, T - \tau_i) - f(B_{\tau_i-}, T - \tau_i)\right]1_{\{\tau_i \le T\}}\right]$$

$$= \mathbb{E}\left[-\int_0^T\left(\sum_{k=1}^K \frac{\lambda_k}{\lambda} M_k^*[f](B_{t-}, T - t)\right)\lambda dt\right.$$

$$\left. + \sum_{i \ge 1}\left[f(B_{\tau_i}, T - \tau_i) - f(B_{\tau_i-}, T - \tau_i)\right]1_{\{\tau_i \le T\}}\right]. \tag{49}$$

Because $M_k^*[f](\cdot, \cdot)$ is continuous on Δ and the budget process is an $\mathbb{F}$-adapted càdlàg process, it follows that $\{M_k^*[f](B_{t-}, T - t)\}_{t \in [0,T]}$ is bounded and $\mathbb{F}$-predictable. Therefore, in terms of the counting process $N = \{N_t\}_{t \ge 0}$ with $N_t = \sum_{i \ge 1} 1_{\{\tau_i \le T\}}$,

for $t \geq 0$, we have

$$
\mathbb{E}\left[\int_0^T \left(\sum_{k=1}^K \frac{\lambda_k}{\lambda} M_k^*[f](B_{t-}, T-t)\right)\lambda dt\right]
$$

$$
= \mathbb{E}\left[\int_{(0,T]} \left(\sum_{k=1}^K \frac{\lambda_k}{\lambda} M_k^*[f](B_{t-}, T-t)\right) dN_t\right]
$$

$$
= \mathbb{E}\left[\sum_{i\geq 1} M_{\kappa_i}^*[f]\left(B_{\tau_i-}, T-\tau_i\right) \cdot 1_{\{\tau_i \leq T\}}\right] \tag{50}
$$

$$
\geq \mathbb{E}\left[\sum_{i\geq 1} M_{\kappa_i}[f]\left(b_i, B_{\tau_i-}, T-\tau_i\right) \cdot 1_{\{\tau_i \leq T\}}\right]
$$

$$
\equiv \mathbb{E}\left[\sum_{i\geq 1} G_{\kappa_i}(b_i)\left(\mu_{\kappa_i} - b_i \rho_{\kappa_i} + \int_0^1 f\left(B_{\tau_i-} - rb_i, T-\tau_i\right) H_{\kappa_i}(dr)\right.\right.
$$

$$
\left.\left. - f\left(B_{\tau_i-}, T-\tau_i\right)\right)1_{\{\tau_i \leq T\}}\right].
$$

Using the conditional distribution of (Z_i, R_i) (with the help of the dominated convergence theorem to interchange the summation and the expectation), we obtain

$$
\mathbb{E}\left[\sum_{i\geq 1}\left(f\left(B_{\tau_i}, T-\tau_i\right) - f\left(B_{\tau_i-}, T-\tau_i\right)\right) 1_{\{\tau_i \leq T\}}\right]
$$

$$
= \mathbb{E}\left[\sum_{i\geq 1} G_{\kappa_i}(b_i)\left(\int_0^1 f\left(B_{\tau_i-} - rb_i, T-\tau_i\right) H_{\kappa_i}(dr) - f\left(B_{\tau_i-}, T-\tau_i\right)\right)1_{\{\tau_i \leq T\}}\right]. \tag{51}
$$

Inserting (50-51) into (49) yields

$$
f(B, T) \geq \mathbb{E}\left[\sum_{i\geq 1} G_{\kappa_i}(b_i)\left(\mu_{\kappa_i} - b_i \rho_{\kappa_i}\right) \cdot 1_{\{\tau_i \leq T\}}\right]
$$

$$
= \mathbb{E}\left[\sum_{i\geq 1}\left(W_i - b_i R_i\right)Z_i \cdot 1_{\{\tau_i \leq T\}}\right], \tag{52}
$$

and this implies that $f(B, T) \geq V(B, T)$ because $(b_i)_{i\geq 1}$ was an arbitrary admissible policy. In particular, with the optimal policy $b^{(\infty)}$ (and its budget process $B^{(\infty)}$) given in Lemma 6, the inequalities in (50) and (52) become equalities, and this proves that $f(B, T) = V(B, T)$. □

References

Abhishek V, Hosanagar K (2013) Optimal bidding in multi-item multislot sponsored search auctions. Oper Res 61(4):855–873

Abrams Z, Mendelevitch O, Tomlin J (2007) Optimal delivery of sponsored search advertisements subject to budget constraints. In: Proceedings of the 8th ACM Conference on Electronic Commerce, EC '07, pp 272–278. ACM

Bae J, Kagel JH (2019) An experimental study of the generalized second price auction. Int J Ind Organ 63:44–68

Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: Approximations and design. Manage Sci 61(4):864–884

Borgs C, Chayes J, Immorlica N, Jain K, Etesami O, Mahdian M (2007) Dynamics of bid optimization in online advertisement auctions. In: Proceedings of the 16th International Conference on World Wide Web, WWW '07, pp 531–540. ACM

Boulatov A, Severinov S (2021) Optimal and efficient mechanisms with asymmetrically budget constrained buyers. Games Econom Behav 127:155–178

Brubach B, Sankararaman KA, Srinivasan A, Xu P (2016) New algorithms, better bounds, and a novel model for online stochastic matching. In ESA

Chaitanya N, Narahari Y (2012) Optimal equilibrium bidding strategies for budget constrained bidders in sponsored search auctions. Oper Res Int J 12(3):317–343

Cholette S, Özlük Ö, Parlar M (2012) Optimal keyword bids in search-based advertising with stochastic advertisement positions. J Optim Theory Appl 152(1):225–244

Dayanik S, Parlar M (2013) Dynamic bidding strategies in search-based advertising. Ann Oper Res 211(1):103–136

Devanur N, Jain K, Sivan B, Wilkens C (2019) Near optimal online algorithms and fast approximation algorithms for resource allocation problems. J ACM 66(1):1–41

Devanur NR, Hayes TP (2009) The adwords problem: online keyword matching with budgeted bidders under random permutations. In: Proceedings of the 10th ACM Conference on Electronic Commerce, EC '09, pp 71–78. ACM

Feldman J, Muthukrishnan S, Pal M, Stein C (2007) Budget optimization in search-based advertising auctions. In: Proceedings of the 8th ACM Conference on Electronic Commerce, EC '07, pp 40–49. ACM

Feldman J, Henzinger M, Korula N, Mirrokni VS, Stein C (2010) Online stochastic packing applied to display ad allocation. In: De Berg M, Meyer U (eds) Algorithms - ESA 2010. Springer, Berlin, pp 182–194

Goel A, Mahdian M, Nazerzadeh H, Saberi A (2010) Advertisement allocation for generalized second-pricing schemes. Oper Res Lett 38(6):571–576

Goel G, Mehta A (2008) Online budgeted matching in random input models with applications to Adwords. In Proceedings of the nineteenth annual ACM-SIAM Symposium on Discrete Algorithms, SODA '08, pp 982–991. SIAM

Hummel P (2018) Hybrid mechanisms for Vickrey–Clarke–Groves and generalized second-price bids. Internat J Game Theory 47(1):331–350

Iyer K, Johari R, Sundararajan M (2014) Mean field equilibria of dynamic auctions with learning. Manage Sci 60(12):2949–2970

Jaillet P, Lu X (2014) Online stochastic matching: New algorithms with better bounds. Math Oper Res 39(3):624–646

Jansen B, Mullen T (2008) Sponsored search: An overview of the concept, history, and technology. Int J Electron Bus 6(2):114–131

Kitts B, Leblanc B (2004) Optimal bidding on keyword auctions. Electron Mark 14(3):186–201

Kotowski MH (2020) First-price auctions with budget constraints. Theor Econ 15(1):199–237

Küçükaydın H, Selçuk B, Özlük Ö (2020) Optimal keyword bidding in search-based advertising with budget constraint and stochastic ad position. J Oper Res Soc 71(4):566–578

Lee D, Zioło P, Han W, Powell WB (2017) Optimal online learning in bidding for sponsored search auctions. In 2017 IEEE Symposium Series on Computational Intelligence (SSCI), pp 1–8

Mahdian M, Nazerzadeh H, Saberi A (2012) Online optimization with uncertain information. ACM Trans Algorith 8(1):1–29

Maillé P, Markakis E, Naldi M, Stamoulis GD, Tuffin B (2012) Sponsored search auctions: An overview of research with emphasis on game theoretic aspects. Electron Commer Res 12(3):265–300

Mehta A, Saberi A, Vazirani U, Vazirani V (2007) AdWords and generalized online matching. J ACM 54(5):22–es

Naor J, Wajc D (2015) Near-optimum online ad allocation for targeted advertising. In Roughgarden T, Feldman M, Schwarz M (eds) Proceedings of the sixteenth ACM Conference on Economics and Computation, EC '15, pp 131–148. ACM

Özlük Ö (2011) Revenue management: A practical pricing perspective, chapter search engine advertising: an overview from a revenue management angle, pp 153–165. Palgrave Macmillan London

Özlük Ö, Cholette S (2007) Allocating expenditures across keywords in search advertising. J Reven Pricing Manage 6(4):347–356

Perlich C, Dalessandro B, Hook R, Stitelman O, Raeder T, Provost F (2012) Bid optimizing and inventory scoring in targeted online advertising. In Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12, pp 804–812. ACM

Pin F, Key P (2011) Stochastic variability in sponsored search auctions: Observations and models. In: Proceedings of the 12th ACM Conference on Electronic Commerce, EC '11, pp 61–70. ACM

Qin T, Chen W, Liu T (2015) Sponsored search auctions: Recent advances and future directions. ACM Trans Intell Syst Technol 5(4):1–34

Rusmevichientong P, Williamson DP (2006) An adaptive algorithm for selecting profitable keywords for search-based advertising services. In: Proceedings of the 7th ACM Conference on Electronic Commerce, EC '06, pp 260–269. ACM

Selçuk B, Özlük Ö (2013) Optimal keyword bidding in search-based advertising with target exposure levels. Eur J Oper Res 226(1):163–172

Skiera B, Abou NN (2013) Prosad: A bidding decision support system for profit optimizing search engine advertising. Market Sci 32(2):213–220