Dynamic Online-Advertising Auctions as Stochastic Scheduling

Ishai Menache and Asuman Ozdaglar
Massachusetts Institute of Technology
{ishai,asuman}@mit.edu

R. Srikant
University of Illinois at Urbana-Champaign
rsrikant@illinois.edu

Daron Acemoglu
Massachusetts Institute of Technology
daron@mit.edu

ABSTRACT

We study dynamic models of online-advertising auctions in the Internet: advertisers compete for space on a web page over multiple time periods, and the web page displays ads in differentiated slots based on their bids and other considerations. The complex interactions between the advertisers and the website (which owns the web page) is modeled as a dynamic game. Our goal is to derive ad-slot placement and pricing strategies which maximize the expected revenue of the website. We show that the problem can be transformed into a scheduling problem familiar to queueing theorists. When only one advertising slot is available on a webpage, we derive the optimal revenue-maximizing solution by making connections to the familiar $c\mu$ rule used in queueing theory. More generally, we show that a $c\mu$-like rule can serve as a good suboptimal solution, while the optimal solution itself may be computed using dynamic programming techniques.

1. INTRODUCTION

Online advertising is a key reason for the continued commercial success of the Internet. A common form of online advertising is the so-called *spot*-auction, which includes adword auctions, banners and other commercial advertising. In spot auctions, conducted by search engines or other web publishers, advertisers bid for space on a webpage. Advertisements are placed either in response to users' web search queries (as in adword auctions [6]), or at predetermined slots on publishers' web pages. Usually, the slots are pre-defined, and there is a clear hierarchy of slots (for example, those near the top of the webpage being more desirable to the advertisers).

While there has been considerable interest in modeling online-advertising auctions, with focus on adword auctions (see, e.g., [5, 15, 1, 9] and [6] for a recent survey), most existing work concentrates on static auction models where a fixed number of advertisers compete for a number of slots during a single time instant. In practice, online advertising auctions are dynamic, with the same (and sometimes changing) set of advertisers competing for slots over time, typically as a function of their history of sales.

In this paper, we study dynamic online advertising auctions. A number of advertisers compete for differentiated slots over multiple time periods. In our baseline model, each advertiser has a single item to sell and the probability of sale depends on the advertiser's type (e.g., quality of product) as well as which advertising slot it has been allocated. Once the single item is sold, the advertiser has no further interest in advertising. We characterize the profit-maximizing auction strategy for a website in this environment.

Following Myerson's seminal contribution [11], we model the characterization of the optimal (profit-maximizing) auction as a mechanism design problem. In particular, because there is commitment on the side of the website to future allocation of slots to advertisers, the revelation principle applies and implies that there is no loss of generality in restricting attention to direct mechanisms in which advertisers report types. Incentive compatibility constraints ensure that no advertiser has an incentive to misreport its type. Moreover, under our baseline assumption that advertiser type is constant over the multiple periods of advertising, a single report by each advertiser at the beginning of the game is sufficient for an appropriately-designed optimal auction. These observations enable a simple mathematical formulation of the optimal dynamic auction and the derivation of analogues of Myerson's characterization results for this dynamic environment.

Our approach not only provides a tractable mathematical formulation of the optimal dynamic auction but it also highlights the parallels between the optimal dynamic auctions and stochastic scheduling problems. In particular, as in stochastic scheduling problems, the website must decide to allocate different advertisers to different slots as a function of their types and the history of success and failure of sales in previous rounds. Importantly, this scheduling must be done in such a way so as to ensure the incentive compatibility of the advertisers (so that they do not misreport their types). Our main results highlight the parallel mathematical structure of dynamic auctions and stochastic scheduling. This parallel structure enables a tight characterization of dynamic optimal auctions in various situations, and in others, it provides us a way of numerically computing the optimal auction.

Related Work. Our paper is related to a number of works in two different areas, which may seem unconnected at first sight. First, our paper is related to the recent literature on dynamic auctions (e.g., [7, 14, 16]). These references consider the problem faced by an auctioneer who sells a finite number of *identical* items to a population of buyers, a framework which is substantially different compared to the multi-period online-advertising domain considered here. Additional recent references related to the online-advertising domain include [12, 4].

Second, the characterization of the optimal allocation rule in our model shares insights and uses tools from the stochastic scheduling and queueing literature ([3], [17], [10], [18] ,[8] and references therein).

Paper structure. Our basic model, consisting of a single advertising-slot, is presented in Section 2. In Section 3, we characterize the optimal auction mechanism for this model. We explicitly derive the optimal allocation mechanism and the corresponding payments in Section 4, under a regularity condition on the valuation distributions. We then proceed to study the multiple advertising slots case. In Section 5, we present the model and characterize the corresponding optimal mechanism. Section 6 restricts attention to the case where the slot qualities obey a certain separability condition, which allows us to employ ideas from stochastic scheduling in order to solve the auction design problem. Conclusions are drawn in Section 7. Due to lack of space, all the proofs of our results are provided in Appendix A.

2. SINGLE SLOT AUCTION

We start with a formal description of the dynamic auction model when a single slot is offered in each time slot over a finite time horizon.

2.1 Model

We assume that there is a single website (or page-owner) who offers a single slot for advertisement. The same slot is offered over M consecutive time slots (or periods). The time slots may correspond, for example, to different periods in a day, or to consecutive days.

The website faces a set of bidders[1] $\mathcal{I} = \{1, \dots, N\}$. Each bidder may use the advertising slot in order to sell a single item. Additional details of the model are provided below.

User Value. Each bidder i is characterized by a scalar quantity t_i which is i's value for selling the item; this value incorporates both the selling price and the scrap value of the item. The uncertainty of the website and all bidders but the ith one about the value t_i is described by a continuous density function f_i over a finite interval $T_i = [a_i, b_i] \subset \mathbb{R}$. The function $F_i : [a_i, b_i] \to [0, 1]$ will denote the cumulative distribution corresponding to the density f_i, so that $F_i(t_i) = \int_{a_i}^{t_i} f_i(s_i)ds_i$. We shall assume that the values of the N bidders are independent. Thus, the joint density function on $T = T_1 \times \cdots \times T_N$ for the vector $t = (t_1, \dots, t_N)$ of the individual values is $f(t) = \prod_{i\in\mathcal{I}} f_i(t_i)$. We shall also denote the joint density of all bidders but the ith one by $f_{-i}(t_{-i}) = \prod_{j\neq i} f_j(t_j)$, where $t_{-i} = (t_1, \dots, t_{i-1}, t_{i+1}, \dots t_n) \in T_{-i}$, $T_{-i} = T_1 \times \cdots T_{i-1} \times T_{i-1} \times \cdots T_N$.

Probability for Selling (PfS). Each bidder i is associated with a probability for selling (PfS), denoted q_i, which is the probability for selling its own item, given that it occupies the slot at any given time m. We assume that the PfS is time-independent. The set $\{q_i\}$ is assumed common knowledge.

[1] In the following, we will use the terms "bidder", "user", and "advertiser" interchangeably.

Dynamic Setup. Once a bidder sells its item, it leaves the auction, and informs the website of its departure. Thus, the participants in each period of the auction are the bidders who had not yet sold their item.

2.2 Dynamic Auction Mechanism

Given the density functions f_i of all users and their associated PfSs $\{q_i\}$, the problem of the website is to select an auction mechanism which would maximize its own expected utility (to be defined below).

An auction mechanism in its full generality can be defined as a pair of allocation and payment functions, as a function of the messages reported by the users. As in Myerson [11], the revelation principle (now for dynamic mechanisms with commitment) implies that attention can be directed without loss of generality to *direct revelation mechanisms*, in which the message set of each user is given by the set of values (or types), i.e., users report their types (correctly or incorrectly) to the website. Since user value is constant over time, a single message at the beginning of the auction provides sufficient information. In view of this property, an auction mechanism is defined formally as follows:

Let $t = (t_1, \dots, t_N)$ denote the vector of types reported by the bidders in the beginning of the auction. Let $A \subset \mathcal{I}$ denote the subset of active bidders (i.e., those who have not sold their item) at some time m (we suppress the dependence of A on time m for notational simplicity). We will sometimes refer to A as the *state* of the auction. The *allocation function (or mechanism)* is an M-tuple $p = (p^1, \dots, p^M)$, where for each m, p^m is a function $p^m : \big(T, A\big) \to [0, 1]^N$ such that if t is the vector of reported types, $p_i^m(t, A)$ is the probability that $i \in \mathcal{I}$ gets the slot at time m. Obviously, $p_i^m(t, A) = 0$ if $i \notin A$. We use the notation $\mathcal{P}$ to denote the set of functions $p = (p^1, \dots, p^M)$. Similarly, the *payment function* is an M-tuple $x = (x^1, \dots, x^M)$, where for each m, x^m is a function $x^m : \big(T, A\big) \to R_+^N$ such that $x_i^m(t, A)$ is the amount of money that user $i \in A$ pays the website for the chance to obtain the slot. Clearly, a user who is not in the system would not be charged, i.e., $x_i^m(t, A) = 0$ for $i \notin A$. We refer to the pair (p, x) as a *(dynamic) auction mechanism*.

2.3 Utilities and Feasible Auction Mechanisms

We next specify the utility functions for the bidders and website. We consider a discounted-cost formulation, with a common discount factor $0 < \delta \leq 1$ for all bidders and the website. Let $U_i(p, x, t_i)$ be the *expected* utility of user i when bidding its true value t_i, assuming that all other users bid their true values. This utility is given by

$$U_i(p, x, t_i) = \int_{T_{-i}} \sum_{m=0}^{M-1} \delta^m \sum_{A\subset\mathcal{I}} P_p^m(A|t)\big(q_i t_i p_i^m(t_i, t_{-i}, A) - x_i^m(t_i, t_{-i}, A)\big) f_{-i}(t_{-i})dt_{-i}, \tag{1}$$

where $P_p^m(A|t)$ is the probability that the state of the auction is A at period m, given the valuation vector t.

The website is interested in his expected profit, given by

$$U_0(p, x, t) = \int_T \sum_{m=0}^{M-1} \delta^m \sum_{A\subset\mathcal{I}} P_p^m(A|t) \sum_{i\in A} \left(x_i^m(t, A)\right) f(t)dt. \tag{2}$$

As in [11], there are three types of constraints that need to be imposed on a mechanism (p, x) for it to be a *feasible* auction mechanism:

1. *Probability constraints (P).*

$$\sum_{i\in A} p_i^m(t,A) \le 1, \quad p_i^m(t,A)\ge 0, \tag{3}$$

for every $m\in[0,M-1]$, $t\in T$, $A\in\mathcal{I}$.

2. *Individual Rationality (IR) constraint.*

$$U_i(p,x,t_i)\ge 0. \quad \text{for every } i\in I, \quad t_i\in T_i. \tag{4}$$

This constraint guarantees that the bidders will participate in the auction.

3. *Incentive compatibility (IC) constraint.* Let t_i be user i's value, and let $s_i\in T_i$ be an alternative value that bidder i may report at the first period. The IC constraint, given by

$$\begin{aligned} U_i(p,x,t_i) \ge \int_{T_{-i}} \sum_{m=0}^{M-1} \delta^m \sum_A P_p^m(A|s_i,t_{-i})\big(q_i t_i p_i^m(s_i,t_{-i},A) \\ - x_i(s_i,t_{-i},A)\big) f_{-i}(t_{-i})dt_{-i}, \end{aligned} \tag{5}$$

ensures that no bidder has any incentive to misreport (or lie about) his value.

The dynamic auction design problem is to choose a mechanism (p,x) so as to maximize $U_0(p,x)$ subject to the P, IR, and IC constraints.

3. CHARACTERIZATION OF THE OPTIMAL MECHANISM

In this section, we show that the dynamic auction design problem can be transformed to another problem, which has a similar structure to Myerson's auction design formulation (see [11]). This structure enables us to obtain explicit characterizations of the optimal allocation and payment functions. Define

$$\hat{p}_i^m(t_i,t_{-i}) \triangleq \sum_A P_p^m(A|t) p_i^m(t_i,t_{-i},A), \tag{6}$$

$$\hat{x}_i^m(t_i,t_{-i}) \triangleq \sum_A P_p^m(A|t) x_i^m(t_i t_{-i},A), \tag{7}$$

and consider the following substitution of variables

$$\tilde{p}_i(t_i,t_{-i}) = q_i \sum_{m=0}^{M-1} \delta^m \hat{p}_i^m(t_i,t_{-i}), \tag{8}$$

$$\tilde{x}_i(t_i,t_{-i}) = \sum_{m=0}^{M-1} \delta^m \hat{x}_i^m(t_i,t_{-i}). \tag{9}$$

The function $\tilde{p}_i(t_i,t_{-i})$ can be viewed as the discounted probability of selling the item, while $\tilde{x}_i(t_i,t_{-i})$ denotes the discounted payment of user i. We refer to the functions $\tilde{p}=(\tilde{p}_1,\dots,\tilde{p}_N)$ and $\tilde{x}=(\tilde{x}_1,\dots,\tilde{x}_N)$ as the *reduced allocation and payment functions* and the pair $(\tilde{p},\tilde{x})$ as the *reduced auction mechanism.*

Using (8)–(9), the utility of the user can be re-written as

$$U_i(p,x,t_i) = \int_{T_{-i}} \left(t_i\tilde{p}_i(t_i,t_{-i}) - \tilde{x}_i(t_i,t_{-i})\right) f_{-i}(t_{-i})dt_{-i} \tag{10}$$

which resembles the one-period utility considered in [11]. The IC constraint (5) obtains the form:

$$U_i(p,x,t_i) \ge \int_{T_{-i}} \left(t_i\tilde{p}_i(s_i,t_{-i}) - \tilde{x}_i(s_i,t_{-i})\right) f_{-i}(t_{-i})dt_{-i}.$$

Finally, the website utility is given by

$$U_0(p,x,t_i) = \int_T \sum_i \tilde{x}_i(t_i,t_{-i}) f(t)dt, \tag{11}$$

which is again similar to the seller's utility in Myerson's work. We define the expected allocation function as $Q_i(\tilde{p},t_i)=\int_{T_{-i}}\tilde{p}_i(t)f_{-i}(t_{-i})dt_{-i}$, where $\tilde{p}_i$ is defined in (8). Hence, $Q_i(\tilde{p},t_i)$ denotes the probability of selling the item when user i reports his value to be t_i, assuming that all other users report their value truthfully. The following definition, related to the expected allocation function, will play a key role in the characterization of the optimal auction mechanism.

Definition 3.1. *A reduced allocation function $\tilde{p}$ is said to be* monotone *if*

$$s_i\le t_i \Rightarrow Q_i(\tilde{p},s_i)\le Q_i(\tilde{p},t_i) \tag{12}$$

for every $i\in\mathcal{I}$ and $s_i,t_i\in T_i$.

The preceding representation of the users and website utilities leads to the following theorem.

Theorem 1. *Consider the following maximization problem*

$$\max_{p\in\mathcal{P}} \int_T \left[\sum_i q_i\left(t_i - \frac{1-F_i(t_i)}{f_i(t_i)}\right)\sum_m \delta^m \hat{p}_i^m(t)\right] f(t)dt \tag{13}$$

subject to the constraints (3),(6) (8), and (12). Let p^* be a solution to this optimization problem, and let $\tilde{p}^*$ be the corresponding reduced allocation function defined in (8). Let

$$\tilde{x}_i^*(t) = \left\{t_i\tilde{p}_i^*(t) - \int_{a_i}^{t_i}\tilde{p}_i^*(s_i,t_{-i})ds_i\right\}, \ \forall i\in\mathcal{I}, t\in T. \tag{14}$$

Then the pair $(\tilde{p}^*,\tilde{x}^*)$ is an optimal reduced auction mechanism.

We emphasize that the optimization problem (13) has a different structure than the problem considered in [11], due to user departures. Nonetheless, as in Myerson's work, (13) can be solved separately for every t, which is a central property that enables us to find the optimal allocation and pricing rule in the dynamic setup considered here. Obtaining an explicit optimal auction mechanism is the subject of the next section.

4. THE REGULAR CASE

In this section, we provide a characterization of the optimal auction mechanism under a similar "regularity" assumption used in [11].

Definition 4.1. *A probability density function $f_i(t_i)$ is said to be regular if $\nu_i(t_i)=t_i-\frac{1-F_i(t_i)}{f_i(t_i)}$ is a monotone strictly increasing function of t_i. The optimal auction design problem is regular if $f_i(t_i)$ is regular for every $i\in\mathcal{I}$.*

We henceforth refer to the function $\nu_i(t_i)$ as the *virtual valuation* of user i.

4.1 Optimal Allocation

We next obtain the optimal allocation function under the regularity assumption. Fix t, and let $\{\nu_i(t_i)\}$ be the virtual valuations of users. With some abuse of notations let $\nu_i=\nu_i(t_i)$.

THEOREM 2. *Enumerate the users such that*

$$q_1\nu_1 \geq q_2\nu_2 \geq \cdots q_N\nu_n, \tag{15}$$

where q_i and ν_i are the selling probability and virtual valuation for user i, respectively. Then the allocation function that maximizes $\sum_{i=1}^{N} q_i\nu_i \sum_{m=0}^{M-1}\delta^m \hat{p}_i^m(t)$ for every $t\in T$ is the one that allocates the slot to users with positive virtual-valuations in increasing user-order (i.e., allocates the slot to user 1 until its item is sold, then to user 2, etc.).

4.2 Optimal Payments

Recall that the optimal reduced payment function can be derived from the optimal reduced allocation according to (14) (see Theorem 1). While (14) determines the reduced (or discounted) payments consistent with an optimal auction, clearly payments per time period are not uniquely determined. This implies that a range of different payment schemes will be consistent with an optimal auction mechanism. Nevertheless, some of the payment schemes are more plausible both on apriori grounds and also because they are easier to implement. In this paper, we present two such alternatives (the second alternative appears in Appendix B due to lack of space).

A central observation that will lead to tractable expressions for the payment functions is the following: The only periods in which user i would have to pay are the periods in which $q_i\nu_i(t_i) \geq q_j\nu_j(t_j)$ for all j currently present in the auction. The actual payment would therefore depend only on the valuations of the other users currently present in the system, and on the *time-to-go*, i.e., the time until the auction ends once user i gains priority.

In view of the above, by specifying the payment function for the user with the highest $q_i\nu_i(t_i)$ at the first slot, we may deduce the pricing formula for any user at any time slot. Indeed, at each time slot we may derive a payment function for the user with the highest $q_i\nu_i(t_i)$ by considering the set of active players and the current time-to-go. To simplify notations, re-enumerate the users according to the terms $q_i\nu_i$ (where $\nu_i \equiv \nu_i(t_i)$), so that $q_1\nu_1 \geq q_2\nu_2 \geq \cdots q_N\nu_n$. We obtain below the payment of user 1, for which the time-to-go is M.

We focus on a specific pricing scheme which ensures incentive compatibility – the *one-shot* payment. This payment is made only at the first period in which the user gets the slot. A user would be willing to make such payment along with a *commitment* of the auctioneer, in the sense that it will keep allocating the advertising slot to the user until it sells its item, or until the auction is over. In order to derive the payment of user 1, we require the following definition.

DEFINITION 4.2. *Consider some user i with positive virtual valuation $\nu_i(t_i)$. A threshold valuation $\bar{t}_i$ is the value that user 1 has to bid so that it has priority over user i, i.e., $q_1\nu_1(\bar{t}_i) = q_i\nu_i(t_i)$, or $\bar{t}_i = \nu_1^{-1}\left(\frac{q_i}{q_1}\nu_i(t_i)\right)$.*

Denote the one-shot payment of user 1 by $x_1^{os}(t)$. Note that $x_1^{os}(t) \equiv \tilde{x}_1^*(t)$, where $\tilde{x}_1^*(t)$ is defined in (14). The next lemma provides a basic characterization of $x_1^{os}(t)$.

LEMMA 1. *Let $\bar{t} = \max\{\bar{t}_2, \nu_1^{-1}(0)\}$. The one-shot payment obeys the following relation*

$$x_1^{os}(t) = q_1\bar{t}\sum_{m=0}^{M-1}\delta^m(1-q_1)^m - q_1\int_{a_1}^{\bar{t}}\sum_{m=0}^{M-1}\delta^m\hat{p}_1^m(s_1,t_{-1})ds_1. \tag{16}$$

Note that if $\bar{t} = \nu_1^{-1}(0)$, (16) immediately implies a payment of $x_1^{os}(t) = q_1\nu_1^{-1}(0)\sum_{m=0}^{M-1}\delta^m(1-q_1)^m$. Assume henceforth that $\bar{t}_2 > \nu_1^{-1}(0)$. This case would result in a lower overall payment compared to the case $\bar{t}_2 = \nu_1^{-1}(0)$, as can be seen from (16). Using the optimal allocation of Theorem 2, we are able to divide the integral in (16) to sub-intervals; each such sub-interval is characterized by inducing a different priority for user 1. For example, assuming that $\nu_3(t_3) > 0$, consider the interval $(\bar{t}_3, \bar{t}_2)$: For each $s_1 \in (\bar{t}_3, \bar{t}_2)$, user 2 has priority over user 1 (while user 1 has priority over the remaining users). Yet, in our multistage auction, user 1 can still obtain the slot, provided that user 2 has sold its item before the auction ends. The potential for obtaining the slot at later periods would result in a discounted payment.

Let $\tilde{M} = \min\{M, |\{i : i > 1 \text{ and } \nu_i(t_i) > 0\}|\}$. The next theorem provides an explicit expression for the one-shot payment.

THEOREM 3. *The optimal one-shot payment is given by*[2]

$$x_1^{os}(t) = q_1\bar{t}\sum_{m=0}^{M-1}\delta^m(1-q_1)^m - \sum_{k=2}^{\tilde{M}}(\bar{t}_k - \bar{t}_{k+1})\prod_{i=1}^{k}q_i\sum_{j=k}^{M}\delta^{j-1}Q(j,k), \tag{17}$$

$$\text{where}\quad Q(j,k) \triangleq \sum_{l_2=0}^{j-k}\sum_{l_3=0}^{j-k-l_2}\cdots \tag{18}$$

$$\cdots\sum_{l_k=0}^{j-k-\sum_{i=2}^{k-1}l_i}(1-q_2)^{l_2}\ldots(1-q_k)^{l_k}(1-q_1)^{j-k-\sum_{i=2}^{k}l_i}$$

and $\bar{t}_{\tilde{M}+1} := \nu_1^{-1}(0)$ if $M > \left|\{i : i > 1 \text{ and } \nu_i(t_i) > 0\}\right|$.

4.3 Multiple Items to Sell

To conclude our treatment for the single-slot case, we consider the scenario in which users may have multiple items to sell during the auction. Due to lack of space, this extension is addressed in Appendix C.

5. THE MULTI-SLOT AUCTION

In this section, we consider the extension of our basic model to multiple advertising slots. We describe the model, and then provide a general characterization theorem of an optimal auction mechanism.

Let $\mathcal{K} = \{1,\ldots,K\}$ be the set of advertising slots. Slots can be of different quality with regard to the selling probabilities. We denote by q_{ik} the probability that user i would sell its item, given that the k-th slot is allocated to it. The probabilities $\{q_{ik}\}$ are assumed time-independent. For simplicity, we shall assume throughout that each bidder has a single item to sell. The assumptions on the user valuations and their distributions remain the same as in Section 2.1.

Whenever possible, we reuse the notations of the single-slot case. The outcome function is now of the form

[2]Empty summations in (17) equal to zero by convention.

$p^m : \big(T, A, m\big) \to [0,1]^{N\times K}$, where $p^m_{ik}(t,A)$ is the probability that $i \in \mathcal{I}$ gets slot k at time m ($p^m_{ik}(t,A) = 0$ if $i \notin A$). The payment functions $x^m : \big(T, A, m\big) \to R_+^{N\times K}$ are such that $x^m_{ik}(t,A)$ is the amount of money which user $i \in A$ pays the website for the chance to obtain the kth slot. Using the above notations, the utility of each user and the utility of the website are given by

$$U_i(p,x,t_i) = \int_{T_{-i}} \sum_{m=0}^{M-1} \delta^m \sum_{A\in\mathcal{I}} P_p^m(A|t) \times$$
$$\times \sum_{k\in\mathcal{K}} \big(t_i q_{ik} p^m_{ik}(t_i,t_{-i},A) - x^m_{ik}(t_i,t_{-i},A)\big) f_{-i}(t_{-i}) dt_{-i},$$

$$U_0(p,x,t) = \int_{T} \sum_{m=0}^{M-1} \delta^m \sum_{A\in\mathcal{I}} P_p^m(A|t) \sum_{i\in A}\sum_{k\in\mathcal{K}} \left(x^m_{ik}(t,A)\right) f(t) dt.$$

The feasibility constraints include (i) the probability constraints, where for every m, t, A

$$\sum_{i\in A} p^m_{ik}(t,A) \; \forall k, \quad \sum_{k\in\mathcal{K}} p^m_{ik}(t,A) \le 1 \; \forall i \in A$$
$$p^m_{ik}(t,A) \ge 0 \quad \forall i \in A, k, \tag{19}$$

(ii) the Individual Rationality (IR) constraint which remains (4), and (iii) the Incentive Compatibility (IC) constraint:

$$U_i(p,x,t_i) \ge \int_{T_{-i}} \sum_{m=0}^{M-1} \delta^m \sum_{A} P_p^m(A|s_i,t_{-i}) \times \tag{20}$$
$$\sum_{k\in\mathcal{K}} \left(q_{ik} t_i p^m_{ik}(s_i,t_{-i},A) - x_{ik}(s_i,t_{-i},A)\right) f_{-i}(t_{-i}) dt_{-i},$$

where $s_i \in T_i$ and t_i is user i's value. As in Section 3, we may express the user utilities in a form that allows us to use Myerson's characterization result of an optimal mechanism. The substitutions of variables are now

$$\hat{p}^m_{ik}(t_i,t_{-i}) \triangleq \sum_A P_p^m(A|t) p^m_{ik}(t_i,t_{-i},A), \tag{21}$$

$$\tilde{p}_i(t_i,t_{-i}) = \sum_{m=0}^{M-1} \delta^m \sum_{k\in\mathcal{K}} q_{ik} \hat{p}^m_{ik}(t_i,t_{-i}), \tag{22}$$

$$\tilde{x}_i(t_i,t_{-i}) = \sum_{m=0}^{M-1} \delta^m \sum_{k\in\mathcal{K}} \hat{x}^m_{ik}(t_i,t_{-i}), \tag{23}$$

where $\hat{x}^m_{ik}(t_i,t_{-i}) \triangleq \sum_A P_p^m(A|t) x^m_{ik}(t_i t_{-i},A)$. Using (22)–(23), the utility of the user, the IC constraint and the website utility are readily seen to be given by (10)–(11) (i.e., the same functions as in the single-slot case). This immediately leads to the following theorem (whose proof follows by by the same arguments as in the proof of Theorem 1, see Appendix A).

THEOREM 4. *Consider the following problem*

$$\max_{p\in\mathcal{P}} \int_T \sum_m \delta^m \left[\sum_i \left(t_i - \frac{1-F_i(t_i)}{f_i(t_i)} \right) \sum_{k\in\mathcal{K}} q_{ik} \hat{p}^m_{ik}(t) \right] f(t) dt \tag{24}$$

subject to the constraints (19), (21), (22) and (12). Let p^* be a solution to this optimization problem, and let $\tilde{p}^*$ be the corresponding reduced allocation function (22). Let

$$\tilde{x}_i^*(t) = \left\{ t_i \tilde{p}_i^*(t) - \int_{a_i}^{t_i} \tilde{p}_i^*(s_i,t_{-i}) ds_i \right\}, \; \forall i \in \mathcal{I}, t \in T. \tag{25}$$

Then $(\tilde{p}^*, \tilde{x}^*)$ is an optimal reduced auction mechanism.

As in the single-slot case, the objective function (24) depends on the allocation only, and furthermore is separable in t. For every t, the maximization problem is a finite horizon dynamic problem, which can be solved numerically, through a backward induction procedure [2]. We emphasize, however, that the solution to (24) should be a monotone allocation in the sense of (12) (due to the incentive compatibility constraint). This constraint essentially couples infinitely-many dynamic programs. Hence, one may consider solving (24) without the constraint, and then verify that (12) holds. In comparison to the single slot case, the monotonicity constraint is generally hard to verify. In the next section, we study a special family of problem instances and consider the validity of the constraint.

6. SEPARABLE SELLING PROBABILITIES

We focus in this section on the case of separable probabilities for selling, analogous to separable click through rates, considered in ad-word auction research (see, e.g., [13]).

6.1 Definitions and Notations

In the separable case, we assume that advertisement slots are arranged in decreasing quality, meaning that the first slot gives the best selling probabilities, uniformly across users. More precisely, each advertising slot k is characterized by some constant $0 < w_k \le 1$, where $w_1 \ge w_2 \cdots \ge w_K$. Each user is associated with a base probability $0 < q_i \le 1$. The PfS of user i at the kth slot is given by $q_{ik} = w_k q_i$ for every $i \in \mathcal{I}$ and $k \in \mathcal{K}$.

We henceforth assume that the valuation distributions $f_i(t_i)$ are regular (see Definition 4.1). Given a valuation vector t (drawn from $f(t)$), the above separability assumption allows us to consider the optimal allocation problem as an equivalent dynamic queuing control problem (or scheduling problem), with the following interpretation: There are N jobs to be preemptively scheduled on the K machines, ordered in decreasing speeds. The reward for processing a job $i \in \mathcal{I}$ is given by the virtual valuation $\nu_i(t_i)$. The objective of the scheduler is to maximize the discounted reward over time.

Based on our experience from the single-slot case, one may conjecture that the optimal allocation is given by the (multi-slot) static $q\nu$ rule, defined: *At each time slot, re-enumerate the present users with positive virtual valuations such that $q_1\nu_1 \ge q_2\nu_2 \ge \cdots$. Then allocate the advertising slots to users with positive virtual-valuations as follows: User 1 gets the first slot, user 2 the second slot, and so on.* It easily follows that the above rule is optimal for a single-period auction. In Appendix D we provide a numerical example which demonstrates that the static $q\nu$ rule *need not* be optimal in a multi-period scenario.

6.2 Optimal Mechanism

As observed in Theorem 4, the optimal allocation problem can be solved separately for every valuation vector t, provided that the resulting allocation is monotone in the sense of (12). We next consider a condition, which is stronger than (12):

$$s_i \le t_i \Rightarrow \tilde{p}_i(s_i,t_{-i}) \le \tilde{p}_i(t_i,t_{-i}) \tag{26}$$

for every $s_i, t_i \in T_i$ and every $t_{-i} \in T_{-i}$. The verification of the monotonicity condition (12), or even of the stronger

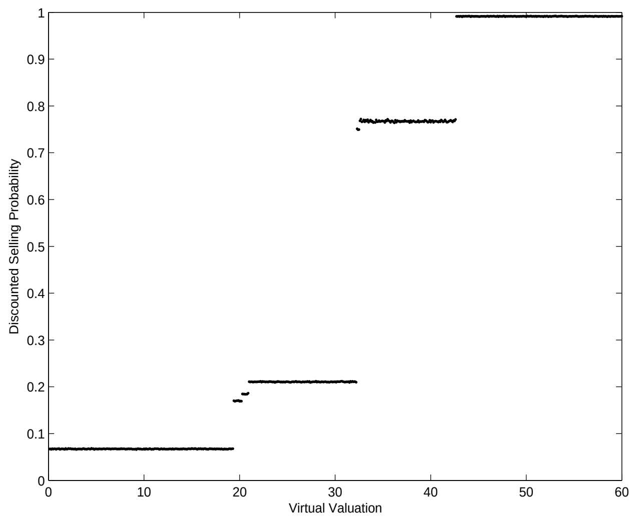

Figure 1: An example for the discounted selling probability of a user as a function of its virtual valuation. In this example there are four users and two non-identical advertisement slots.

condition (26) turns out to be a difficult task, which we have not resolved analytically. To the best of our knowledge, there are no tools in the dynamic programming literature for examining such monotonicity conditions on a particular subset of the program's "rewards". Yet, we have examined numerically whether (26) holds, and in all our experiments, (26) was indeed satisfied.

Figure 1 demonstrates a typical simulation instance of a two slot auction with 4 users. The figure shows that the reduced allocation function is monotonously increasing in the virtual valuation (hence in the valuation itself), as required in (26). Furthermore, the relatively long horizontal lines correspond to subsets of values in which the allocation function for the particular user is fixed, according to some static priority. We observe some fluctuations between two solid lines, which reflect dynamic and perhaps non-stationary allocation rules (as demonstrated in the numerical example of Appendix D).

The implications of the above observations are twofold. On the positive side, the empirical verification of the stronger monotonicity condition (26) suggests that even if cases where (26) does not hold were to exist, then due to their rareness, they would be averaged out while considering the original monotonicity constraint (12). On the negative side, we point to the relation between Fig. 1 and the price setting procedure. Observe that the integral in (25) is equivalent to the area below the graph curve, where the x-axis range is limited to $[0, \nu_i]$. Due to the transient behavior between solid lines, a fine quantization of the value spaces might be required in order to obtain the payments in an adequate precision.

6.3 The $q\nu$-Based Mechanism

We next consider an allocation mechanism which is based on the the $q\nu$ rule (defined in Section 6.1), which enables an accurate and easier payment determination, compared to the optimal allocation rule. The next theorem establishes the incentive compatibility of the $q\nu$ allocation.

THEOREM 5. *The $q\nu$ allocation is a monotone allocation.*

Denote by $p^{q\nu}$ the $q\nu$-based allocation. An immediate consequence of the above theorem is that we may use (25) in order to obtain the one-shot incentive compatible payments. This property is summarized in the next proposition.

PROPOSITION 6. *Let $\tilde{p}_i^{q\nu}$ and $\tilde{x}_i^{q\nu}$ be the reduced allocation and payment functions of user i under the allocation $p^{q\nu}$. Then the relation between $\tilde{p}_i^{q\nu}$ and $\tilde{x}_i^{q\nu}$ which ensures incentive compatibility is obtained through (25), with $\tilde{p}_i^{q\nu}$ and $\tilde{x}_i^{q\nu}$ substituting $\tilde{p}_i^*$ and $\tilde{x}_i^*$ respectively.*

Observe from (25) that the basic quantity that is required for price derivation is $\tilde{p}_i$, the discounted selling probability. Unlike the single slot case, obtaining a closed-form solution for $\tilde{p}_i$ becomes an intractable task, even for the more simple $q\nu$ rule: Once the auction goes beyond two periods, $\tilde{p}_i$ has to be estimated through Monte-Carlo like experiments. Even with $\tilde{p}_i$ at hand, the integral in (25) induces an additional difficulty for obtaining the user payments, as discussed in Section 6.2. In this context, the $q\nu$-based mechanism has a fundamental advantage over the optimal mechanism. The $q\nu$-based mechanism allows one to consider a *finite* number of sub-intervals in which $\tilde{p}_i$ remain constant. The number of these intervals is upper-bounded by the number of users (as in the single-slot case).

The $q\nu$-based mechanism thus provides a much simpler alternative for the auction problem, both in terms of the easier-to-implement static allocation rule, and also with regard to the payments derivation. In order to further motivate its use, we examine through simulations the revenue gap between the optimal mechanism and the $q\nu$ mechanism. Our simulation results indicate that the average performance gap between the two mechanisms is not greater than 3.2% (see Appendix E for more details on the performed experiments). To supplement our experimental results, we refer the reader to the literature in stochastic scheduling, which suggests that the $q\nu$ rule (or the $c\mu$ rule when referring to holding costs) provides an excellent heuristic as the number of users (or job) grows (see [18] for a survey).

7. CONCLUSION

In this paper, we studied dynamic online-advertising auctions. We provided mathematically tractable models of the dynamic auction design problem, where the goal is to maximize the expected profit of the website. We established connections between dynamic auction design problems and stochastic scheduling problems that are common in the queueing theory literature. By exploiting this connection, we showed that $c\mu$-like allocation rules are either optimal or near-optimal, depending on the model studied.

We briefly note several extensions and interesting open directions: The framework that we use here is a Bayesian auction framework, where it is assumed that the auctioneer has some knowledge of the firms' values, and moreover perfect knowledge of the selling probabilities. It is of great interest to consider the case where such knowledge is not available, which may require novel methods and techniques to solve the mechanism design problem. Another challenging venue for future work is to examine the case where the selling probabilities vary over time. In this case, the $c\mu$ rule need not be optimal, even for the single-slot case. Consequently, the allocation mechanism itself would require algorithmic solutions that have not been examined within the queuing research community.

8. REFERENCES

[1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *ACM conference on Electronic commerce*, pages 1–7, 2006.

[2] D. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 1995.

[3] C. Buyukkoc, P. Varaiya, and J. Walrand. The $c\mu$ rule revisited. *Advances in Applied Probability*, 17(1):237–238, 1985.

[4] F. Constantin, J. Feldman, S. Muthukrishnan, and Martin Pal. An online mechanism for ad slot reservations with cancellations. In *SODA*, 2009.

[5] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[6] J. Feldman and S. Muthukrishnan. Algorithmic methods for sponsored search advertising. In *Sigmetrics (Tutorial Sessions)*, pages 91–124, 2008.

[7] J. Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54(2):291–310, 2006.

[8] K. D. Glazebrook and J. Nino-Mora. Parallel scheduling of multiclass M/M/m queues: Approximate and heavy-traffic optimization of achievable performance. *Operations Res.*, 49(4):609–623, 2001.

[9] G. Iyengar and A. Kumar. Characterizing optimal adword auctions. Technical Report CORC TR-2006-04, Columbia University, November 2006.

[10] T. Kampke. On the optimality of static priority policies in stochastic scheduling on parallel machines. *Journal of Applied Probability*, 24(2):430–448, 1987.

[11] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[12] H. Nazerzadeh, A. Saberi, and R. Vohra. Dynamic cost-per-action mechanisms and applications to online advertising. In *WWW*, pages 179–187, 2008.

[13] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK, 2007.

[14] M. Pai and R. Vohra. Optimal dynamic auctions. Technical Report 1461, Northwestern University, 2008.

[15] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.

[16] G. Vulcano, G. Van-Ryzin, and C. Maglaras. Optimal dynamic auctions for revenue management. *Management Science*, 48(11):1388–1407, 2002.

[17] J. Walrand. *An introduction to Queuing Networks*. Prentice Hall, New Jersey, 1988.

[18] G. Weiss. A tutorial in stochastic scheduling. *Scheduling Theory and Its Applications*, 1995.

APPENDIX

A. PROOFS

Proof of Theorem 1. (outline) Once expressing the bidder and website utilities as (10) and (11), the proof follows straightforwardly from [11]. We follow the lines of Lemmas 2 and 3 in [11] and obtain (13), which is identical to the objective function in Lemma 3 of [11] when applying the transformation $\tilde{p}_i(t_i, t_{-i}) = q_i \sum_{m=0}^{M-1} \delta^m \hat{p}_i^m(t_i, t_{-i})$. In particular, these lemmas shows that the IC constraint (5) is equivalent to (12), while the IR constraint (4) is incorporated in the optimal reduced payment function (14). It should be noted that Myerson's derivation does not require any constraints on p (constraints are imposed later on, when deriving explicit solutions), which allows us to use our transformations without considering the specific constraints (6) and (8). □

Proof of Theorem 2. Note first that the allocation mechanism (15) is monotone (see Definition 3.1), as due to regularity, when a user raises its bid, it may only improve its priority in the sense of (15), hence increase its allocation probability. The maximization of $\sum_{i=1}^{N} q_i \nu_i \sum_{m=0}^{M-1} \delta^m \hat{p}_i^m(t)$ can be viewed in the context of dynamic control of queuing systems, where the advertising slot is regarded as a single server, which serves user i at a rate proportional to q_i. We use below the proof idea of [3].

The proof follows by induction on the number of periods M. Obviously, for $M = 1$ the assertion is true. Suppose the assertion is true for some M and consider an auction with $M + 1$ periods. By induction, we assume that the optimal mechanism follows (15) for $m = 1, \ldots, M$. By way of contradiction, suppose that the optimal mechanism does not follow (15) for the first period $m = 0$. Then there are users j and i at the first period such that $q_i \nu_i > q_j \nu_j$ (i.e., $i < j$), while the optimal mechanism assigns the slot to user j. Denote by τ the *first* period that the mechanism assigns the slot to user i (set $\tau = 0$ if user i is not assigned the slot). Note that τ is a random variable. Since the optimal mechanism follows (15) for $m \geq 1$, it does not serve j during periods $m = 1, \ldots, \tau$. We next show that this mechanism can be improved upon, by interchanging the actions at times 0 and τ.

We modify the above mechanism by (i) assigning the slot to user i at time $m = 0$ (ii) assigning the slot to user j at time τ. Recall that this assignment is valid since user j is in the system at time τ, as users are served according to (15) for $m \geq 1$. Observe that the states under both policies will be the same from time $\tau + 1$ onwards. Hence, the effect of this modification can be calculated as follows. For each sample-path of the original mechanism that serves user j at time 0 while user i is still in the system, the surplus from user i increases by $q_i \nu_i (1 - \delta^\tau)$, whereas the benefit from user j decreases by $q_j \nu_j (1 - \delta^\tau)$. Thus, the expected change in the mechanism's value is $(q_i \nu_i - q_j \nu_j)\mathbb{E}(1 - \delta^\tau)$, which is positive since $q_i \nu_i > q_j \nu_j$. To conclude the proof, we introduce a user with $q_0 = \nu_0 = 0$, whose inclusion yields that not allocating the slot while there are users with positive virtual valuations is a suboptimal action. This shows that (15) is optimal. □

Proof of Lemma 1. Using (14) and recalling that $x_1^{os}(t) \equiv \tilde{x}_1(t)$ we have

$$x_1^{os}(t) = q_1 t_1 \sum_{m=0}^{M-1} \delta^m \hat{p}_1^m(t) - q_1 \int_{\bar{t}}^{t_1} \sum_{m=0}^{M-1} \delta^m \hat{p}_1^m(s_1, t_{-1}) ds_1 - q_1 \int_{a_1}^{\bar{t}} \sum_{m=0}^{M-1} \delta^m \hat{p}_1^m(s_1, t_{-1}) ds_1. \quad (27)$$

Noting that user 1 gets the slot with probability one as long as its bid is above $\bar{t}$ (and as long as user 1 is in the system), (27) is equivalent to

$$x_1^{os}(t) = q_1 t_1 \sum_{m=0}^{M-1} \delta^m (1-q_1)^m - q_1(t_1 - \bar{t}) \sum_{m=0}^{M-1} \delta^m (1-q_1)^m$$
$$- q_1 \int_{a_1}^{\bar{t}} \sum_{m=0}^{M-1} \delta^m \hat{p}_1^m(s_1, t_{-1}) ds_1,$$

which leads to (16). □

Proof of Theorem 3. The proof follows by a couple of combinatorial arguments. For each value in the interval $(\bar{t}_k - \bar{t}_{k+1})$, user 1 may be able to sell its item as of the kth period; user 1 would advertise its item only if all other users $i = 2, \ldots, k$ sold their items; this explains the term $\prod_{i=1}^{k} q_i$ in (17). In each of the periods $j \in [k, M]$, user 1 may sell its item with positive probability $Q(j,k) \prod_{i=1}^{k} q_i$, where the benefit from selling is discounted by a multiplicative term δ^{j-1}. The term $Q(j,k)$ stands for the probability of having exactly $j-k$ sale failures out of j attempts, which eventually result in user 1 selling the item at the jth period. □

Proof of Theorem 5. The main idea behind the proof is to apply a sample-path argument, which suggests that for every period m, if the user is still in the system, then its probability for selling the item increases with its bid. This property is summarized below.

LEMMA 2. *Denote by $\pi_i^m(t_i)$ the probability that user i sells its item in the mth period, given that it has not sold it up to that period. Then $s_i \le t_i \Rightarrow \pi_i^m(s_i) \le \pi_i^m(t_i)$.*

PROOF. Fix t_{-i} and consider two user valuations s_i, t_i such that $s_i \le t_i$. Let $w_i^m(t_i)$ be the weight of the slot assigned to user i at period m as a function of its valuation t_i ($w_i^m(t_i) \equiv 0$ if no slot is assigned). Note that $w_i^m(t_i)$ is a random variable. Observe that the claim trivially holds for $m = 0$, as the higher the value, the (weakly) better is the obtained slot, and consequently $\pi_i^m(s_i) = q_i w_i^m(s_i) \le q_i w_i^m(t_i)$. For every $m > 0$, consider any sample path of user departures. Specifically, focus on the departure sample-path of users with valuations $\nu_j(t_j) > \nu_i(t_i) \ge \nu_i(s_i)$. Note that each such sample-path occurs with the same probability regardless if user i valuation is s_i or t_i. For each sample path, denote by $\Delta_i^m(t_i)$ the number of departures of users j with $\nu_j(t_j) > \nu_i(t_i) \ge \nu_i(s_i)$ that result in an upgrade in user i's slot at time m, given that the user valuation is t_i. Obviously, $\Delta_i^m(t_i) \ge \Delta_i^m(s_i)$. It immediately follows from the last assertion that

$$w_i^m(s_i) \le w_i^m(t_i). \tag{28}$$

Indeed, let $b_i^0(t_i)$ be the index of the server allocated to user i at period m as a function of t_i. Then

$$b_i^m(s_i) \ge b_i^0(s_i) - \left[b_i^0(s_i) - b_i^0(t_i)\right] - \Delta_i^m(s_i)$$
$$\ge b_i^0(s_i) - \Delta_i^m(t_i) = b_i^m(t_i), \tag{29}$$

where the first inequality follows from the optimistic scenario that all users j with virtual valuations $\nu_i(s_i) \le \nu_j(t_j) \le \nu_i(t_i)$ sell their items by time m. Inequality (29) immediately implies (28) by the separability assumption. □

We are now ready to prove the theorem. To that end, we reuse the notation $\pi_i^m(t_i)$ the probability that user i sells its item in the mth period, given that it has not sold it up to that period. We establish the stronger monotonicity property (26). Fixing t_{-i}, it can be easily seen that $\tilde{p}_i(t_i) = \pi_i^0(t_i) + \delta(1-\pi_i^0(t_i))\pi_i^1(t_i) + \cdots + \delta^{M-1} \prod_{m=0}^{M-2}(1-\pi_i^m(t_i))\pi_i^{M-1}(t_i)$. We use the last relation in order to prove (26). The key property for the proof is $\pi_i^m(s_i) \le \pi_i^m(t_i)$, obtained in Lemma 2.

We prove a more general statement: Let $\tilde{\pi} = (\tilde{\pi}^0, \ldots, \tilde{\pi}^{M-1})$ and $\pi = (\pi^0, \ldots, \pi^{M-1})$ be any two M-dimensional probability vectors such that $\tilde{\pi} \le \pi$ (component-wise). Then,

$$\tilde{\pi}^0 + \delta(1-\tilde{\pi}^0)\tilde{\pi}^1 + \cdots + \delta^{M-1}\tilde{\pi}^{M-1} \prod_{m=0}^{M-2}(1-\tilde{\pi}^m) \tag{30}$$
$$\le \pi^0 + \delta(1-\pi^0)\pi^1 + \cdots + \delta^{M-1}\pi^{M-1} \prod_{m=0}^{M-2}(1-\pi^m).$$

The proof of the above statement follows by induction on M. The statement clearly follows for $M = 0$. Assume it holds for M-dimensional vectors, and consider two $M+1$ probability vectors $\tilde{\pi} = (\tilde{\pi}^0, \ldots, \tilde{\pi}^M)$ and $\pi = (\pi^0, \ldots, \pi^M)$ such that $\tilde{\pi} \le \pi$. Then

$$\tilde{\pi}^0 + \delta(1-\tilde{\pi}^0)\pi^1 + \cdots + \delta^M \tilde{\pi}^M \prod_{m=0}^{M-1}(1-\tilde{\pi}^m)$$
$$\le \tilde{\pi}^0 + (1-\tilde{\pi}^0)\delta\left(\pi^1 + \delta(1-\pi^1)\pi^2 + \cdots + \delta^{M-1}\Pi^m\right)$$
$$\le \pi^0 \cdot 1 + (1-\pi^0)\delta\left(\pi^1 + \delta(1-\pi^1)\pi^2 + \cdots + \delta^{M-1}\Pi^m\right),$$

where $\Pi^M \triangleq \tilde{\pi}^M \prod_{m=1}^{M-1}(1-\tilde{\pi}^m)$; the first inequality follows from the induction hypothesis, and the second inequality since $\delta\big(\pi^1 + \delta(1-\pi^1)\pi^2 + \cdots + \delta^{M-1}\pi^M \prod_{m=1}^{M-1}(1-\pi^m)\big) \le 1$ The above yields that (30) holds for every M, and thus (26) holds as a special case. □

Proof of Proposition 6. For the proof, we use certain properties that were obtained in Lemmas 2-3 in [11]. Their use is justifiable by the substitution of variables (22) and (23) that leads to user utilities (10) and website utility (11), which have the same form as in [11].

Lemma 2 in [11] gives an equivalent characterization of a feasible mechanism. Namely, a mechanism is feasible if and only if the probability constraints (19) and the monotonicity constraint (12) hold, along with the two following conditions:

$$U_i(p, x, t_i) = U_i(p, x, a_i) + \int_{a_i}^{t_i} Q_i(\tilde{p}, s_i) ds_i, \tag{31}$$

for every $i \in \mathcal{I}$ and every $t_i \in [a_i, b_i]$, and

$$U_i(p, x, a_i) \ge 0, \quad i \in \mathcal{I}. \tag{32}$$

Using the analysis of Lemma 3 in [11], the website problem is to maximize

$$\max_{p \in \mathcal{P}} \int_T \left[\sum_i \left(t_i - \frac{1-F_i(t_i)}{f_i(t_i)}\right)\tilde{p}_i(t)\right] f(t)dt - \sum_i U_i(p, x, a_i), \tag{33}$$

subject to (19), (12), (31) and (32). In this formulation, x appears only in the lest term of the objective function and in the constraints (31)–(32). These two constraints may be written as

$$\int_{T_{-i}} \left\{ t_i \tilde{p}_i(t) - \int_{a_i}^{t_i} \tilde{p}_i(s_i, t_{-i}) ds_i - \tilde{x}_i(t) \right\} f_{-i}(t_{-i}) dt_{-i} = U_i(p, x, a_i) \geq 0, \quad (34)$$

for every $i \in \mathcal{I}$ and $t_i \in [a_i, b_i]$. If the website fixes its allocation mechanism to the $q\nu$ rule, then (19) and (12) hold (by Theorem 5). The utility (33) is not maximized, however *given* the allocation mechanism, if $\tilde{x}^{q\nu}$ is chosen according to (25), then it satisfies both (31) and (32) and it obtains $\sum_i U_i(p, x, a_i) = 0$, which is the best possible value for this term in (33). □

B. MULTI-PERIOD PAYMENTS

The objective of this section is to obtain a set of payments that are paid at each period in which the user is still in the system. In general, there are several ways to design such payments. We shall focus, however, on a specific implementation that has a recursive structure. We note that the pricing scheme we suggest here is theoretically equivalent to the one-shot payment scheme for all parties involved, yet could be psychologically more plausible for the users to pay per period instead of transferring all money at once.

Recall that $x_1^{os}(t) = \tilde{x}_1^*(t) = \sum_{m=0}^{M-1} \delta^m \hat{x}_1^m(t)$ and that $\tilde{x}_1^*(t)$ can be regarded as the average (discounted) payment. The user will be present in the system at period $m \in [0, M-1]$ with probability $(1-q_1)^m$, hence $x_1^{os}(t) = \tilde{x}_1(t) = \sum_{m=0}^{M-1} \delta^m (1-q_1)^m x_1^m(t)$, where $x_1^m(t)$ stands for the possible payment at period m. The technique for constructing the per-period payments can be roughly described as follows. The payment at a given period m is comprised of all summands at the right-hand-side of (17) which originally include the term $(1-q_1)^m$. The term $(1-q_1)^m$ itself would not be included in the payment, and moreover all terms have to be divided by δ^m to account for the delayed payment. This gives rise to the following theorem (in which we use Definition 4.2 for a threshold valuation $\bar{t}_k$).

THEOREM 7. *An optimal payment mechanism is given by*

$$x_1^h(t) = q_1 \bar{t} - \sum_{k=2}^{\tilde{M}} (\bar{t}_k - \bar{t}_{k+1}) \prod_{i=1}^{k} q_i \sum_{j=k}^{M} \delta^{j-h-1} \bar{Q}(j-h, k) \quad (35)$$

for every $h \in [0, M-1]$, where for every $J \geq k$,

$$\bar{Q}(J,k) \triangleq \sum_{l_2=0}^{J-k} \sum_{l_3=0}^{J-k-l_2} \cdots \quad (36)$$
$$\cdots \sum_{l_{k-1}=0}^{J-k-\sum_{i=2}^{k-2} l_i} (1-q_2)^{l_2} \cdot \cdot (1-q_{k-1})^{l_{k-1}} (1-q_k)^{J-k-\sum_{i=2}^{k-1} l_i}$$

($\bar{Q}(J,k) = 0$ for $J < k$), and $\bar{t}_{\tilde{M}+1}$ is set to $\nu_1^{-1}(0)$ in case that $M > \left|\{i : i > 1 \text{ and } \nu_i(t_i) > 0\}\right|$.

PROOF. The proof idea is to show that $\sum_{m=0}^{M-1} \delta^m (1-q_1)^m x_1^m(t)$ and the one-shot payment $x_1^{os}(t) = \tilde{x}_1^*(t)$ are equivalent. To that end, we show that the payment at slot $h \in [0, M-1]$ includes all summands in (17) with the term $(1-q_1)^h$, after dividing those by $(1-q_1)^h \delta^h$.

The one-shot payment can be written as

$$x_1^{os}(t) = \tilde{x}_1^*(t) = q_1 \bar{t} \sum_{m=0}^{M-1} \delta^m (1-q_1)^m \quad (37)$$
$$- \sum_{k=2}^{\tilde{M}} (\bar{t}_k - \bar{t}_{k+1}) \prod_{i=1}^{k} q_i \sum_{j=k}^{M} \delta^{j-1} Q(j,k) =$$
$$= q_1 \bar{t} \sum_{m=0}^{M-1} \delta^m (1-q_1)^m$$
$$- \sum_{k=2}^{\tilde{M}} (\bar{t}_k - \bar{t}_{k+1}) \prod_{i=1}^{k} q_i \sum_{j=k}^{M} \delta^{j-1} \sum_{h=0}^{j-k} (1-q_1)^h \bar{Q}(j-h, k).$$

Summing all terms above that contain $(1-q_1)^h$, and then dividing by $(1-q_1)^h \delta^h$, immediately leads to (35). □

Fixing the bid vector t, the multi-period payments derived in (35) are seen to be a function of the period m, a property that is appealing from a practical point of view. In our context, if the auction was to start at some period $m' > 1$, the payments for user 1 would be the same as in the case where the auction started at $m = 0$ and the user has not sold its item up to period m'. Put differently, the set of payments from each $m' > 1$ are optimal and incentive compatible, even when ignoring the past $m = 0, \ldots, m' - 1$.

C. SINGLE SLOT WITH MULTIPLE ITEMS PER USER

We consider an extension to the single-slot model, in which users may have multiple items to sell during the auction. As a concrete model for this extension, we study here the case where each user i has a fixed number of items J_i to sell, and it remains in the auction as long as its items were not sold. The J_i's are assumed common knowledge (e.g., declared by the bidder when joining the auction). Obviously, one may think of other models that incorporate multiple sales per-user (e.g., models that include probabilistic assumptions on the number of items to be sold by each user), however they are beyond the scope of the current paper.

Assume that the user's valuation t_i and the selling probability q_i remain fixed and do not depend on the number of items sold. Under this assumption, it can be easily seen that the user and website utilities can be written as in (1) and (2) respectively, with the auction state A being the subset of active users along with the number of items that each one of them has sold. Being able to represent the utilities in the same manner as before, we may use the transformations (8) and (9). These transformations will lead to a result equivalent to Theorem 1, and, again, turn the auction design problem into a stochastic scheduling problem.

The most significant property that allows us to solve the auction problem for the case of "rejoining" users, is that the optimal allocation mechanism continues to be the static $q\nu$ rule, which was proven to be optimal in Theorem 2. As a matter of fact, the $q\nu$ rule holds in single-server queuing systems under fairly general assumptions, which accommodate in particular user (or job) arrivals, as long as the user types (q_i, ν_i) are fixed throughout (see, e.g., [3])[3].

[3]Accordingly, when considering other models that allow for multiple sells per user, if the webpage utility can be written as (13), then the $q\nu$ rule remains optimal. This implies that

Based on the optimality of the $q\nu$ rule, we are able to derive explicit pricing formulae. For simplicity, we focus on one-shot payment, and obtain an explicit pricing formula in the spirit of Theorem 3. *Whenever possible, we use the notations and definitions of Section 4.*

Denote by $\acute{p}_i(M')$ the discounted average number of items that user i sells, given that it gets the advertising slot for M' consecutive periods (by "discounted" we mean that a sale at the mth period is multiplied by δ^{m-1}). Formally, $\acute{p}_i(M') = \mathbb{E}\left\{\sum_{m=0}^{\tilde{M}-1} \delta^m 1\{\text{an item is sold at time } m\}\right\}$. The lemma below derives an explicit expression for $\acute{p}_i(M')$.

LEMMA 3. *The discounted average number of items that user i sells, given that it gets the advertising slot for M' periods, is given by*

$$\acute{p}_i(M') = q \sum_{m=0}^{M'-1} \delta^m \sum_{j=0}^{J_i-1} q^j (1-q)^{m-1-j} \binom{m-1}{j}. \quad (38)$$

PROOF. The proof follows from the following equation

$$\acute{p}_i(M') = \mathbb{E}\left\{\sum_{m=0}^{\tilde{M}-1} \delta^m 1\{\text{an item is sold at time } m\}\right\}$$

$$= \sum_{m=0}^{M'-1} \delta^m \mathbb{P}\{\text{an item is sold at time } m\}$$

$$= q \sum_{m=0}^{M'-1} \delta^m \mathbb{P}\{\text{\# of items sold by time } (m-1) \text{ is } \leq J_i - 1\},$$

which is equivalent to (38). □

Having $\acute{p}_i(M')$ at hand, we proceed to obtain the pricing formula for $i = 1$, assuming that the time horizon is M periods. Define $S_k = \sum_{i=2}^{k} J_i$, and let $\bar{N} = |\{i : \nu_i(t_i) > 0\}|$. We further define the variable $\tilde{K}$ as follows: If there exists a user index $\bar{i}$ such that (i) $S_{\bar{i}} < M$ and (ii) $S_{\bar{i}+1} \geq M$ and (iii) $2 \leq \bar{i} \leq \bar{N}$, set $\tilde{K} = \bar{i}$, otherwise set $\tilde{K} = \bar{N}$. With these definitions, we have the following price characterization.

THEOREM 8. *The optimal one-shot payment is given by*

$$x_1^{os}(t) = \bar{t}\acute{p}_1(M) - \sum_{k=2}^{\tilde{K}} (\bar{t}_k - \bar{t}_{k+1}) \prod_{i=2}^{k} q_i^{J_i} \sum_{j=S_k+1}^{M} \delta^{j-1} \bar{Q}(j, S_k) \acute{p}_1(M-j), \quad (39)$$

where $\bar{Q}(j, S_k)$ is given by (36), $\acute{p}_i(M)$ by (38), and $\bar{t}_{\tilde{K}+1} := \nu_1^{-1}(0)$ if $\tilde{K} = \bar{N}$.

PROOF. (outline) Follows similarly to the proof of Theorem 3. The main differences from that proof are (i) User i will leave the system only if it has J_i successes, hence the term $\prod_{i=2}^{k} q_i^{J_i}$. (ii) The formula includes the function $\acute{p}_i(M)$, to account for multiple sales of the user. □

even if prices cannot be computed explicitly, they can at least be obtained numerically.

D. SUBOPTIMALITY OF THE Qν RULE FOR THE MULTI-SLOT CASE

We show through a numeric example that the static $q\nu$ rule need not be optimal when considering multiple advertising slot auctions. The example is constructed in the spirit of [10]. Consider a ten-period auction of two advertising slots with identical quality $w_1 = w_2 = 1$. The auction consists of four users with virtual valuations: $\nu_1 = 100$, $\nu_2 = 100$, $\nu_3 = 2$ $\nu_4 = 15$, and base selling probabilities $q_1 = 0.8100$, $q_2 = 0.0081$, $q_3 = 0.081$, $q_4 = 0.0103$. Note that $q_1\nu_1 > q_2\nu_2 > q_3\nu_3 > q_4\nu_4$. The discount factor δ is set to 0.8. The allocation mechanism that maximizes (24) is given by the following *dynamic priority* rule: (1) Start the auction by allocating the two slots to users 1 and 2; (2) If user 1 sells its item before user 2, continue according to the priority order $2 \to 3 \to 4$ (i.e., user 3 gets a slot after user 1 departs); (3) If user 2 sells its item before user 1, continue according to the priority order $1 \to 4 \to 3$ (i.e., user 4 gets a slot after user 2 departs). An exception to this rule is the last period, in which if users 1, 3 and 4 are in the system, then 1 and 3 get the slots. Obviously, this exception follows from the fact that the $q\nu$ rule is optimal for the last period. Fixing all problem parameters and decreasing just q_4 to 0.009 leaves rules (1)–(2) above the same; rule (3) remains valid until the 5th period, from which the optimal priority order is $1 \to 3 \to 4$ (meaning that optimal allocation might be non-stationary).

E. OPTIMAL REVENUES VS. $Q\nu$-BASED REVENUES

We have compared through simulations the expected revenues in the optimal mechanism with the revenues obtained in the $q\nu$ mechanism. We briefly describe here the simulation setup and methodology. Recall that the expected revenue of the website is given by $U_0(p, x, t_i) = \int_T \sum_i \tilde{x}_i(t_i, t_{-i}) f(t) dt$, where $\tilde{x}_i(t)$ is the expected payment of user i for a given valuation vector t. Hence, in order to estimate $U_0(p, x, t_i)$, a significant number of bid vectors is required to be generated according to the underlying distribution $f(t)$. For each generated vector, the total payments of the users can be obtained numerically according to the reduced payment function (25), which holds for both mechanisms. The overall simulation process for estimating revenues is thus time consuming and obviously grows with the number of users. In order to provide relatively accurate estimates of the revenue gap under both mechanisms, we consider in our experiments relatively small networks (up to 8 users and 3 slots) with a number of periods $M \leq 10$. The user base probabilities q_i and the slot qualities w_k are drawn at random. We further assume that the user valuations are drawn uniformly between $[0, 100]$. For each set of $\{q_i\}$ and $\{w_k\}$, we draw a large number of valuation vectors and average the revenues for each mechanism.

Our simulation results indicate that the average performance gap between the optimal mechanism and the $q\nu$ mechanism is approximately 1.5%, where we did not observe a gap greater than 3.2% in any simulation instance.