

Efficient Advert Assignment

Frank Kelly¹, Peter Key², and Neil Walton³

¹University of Cambridge

²Microsoft Research Cambridge

³University of Amsterdam

Abstract

We develop a framework for the analysis of large-scale Ad-auctions where adverts are assigned over a continuum of search types. For this pay-per-click market, we provide an efficient mechanism that maximizes social welfare. In particular, we show that the social welfare optimization can be solved in separate optimizations conducted on the time-scales relevant to the search platform and advertisers. Here, on each search occurrence, the platform solves an assignment problem and, on a slower time-scale, each advertiser submits a bid which matches its demand for click-throughs with supply. Importantly, knowledge of global parameters, such as the distribution of search terms, is not required when separating the problem in this way. Exploiting the information asymmetry between the platform and advertiser, we describe a simple mechanism which incentivizes truthful bidding and has a unique Nash equilibrium that is socially optimal, and thus implements our decomposition. Further, we consider models where advertisers adapt their bids smoothly over time, and prove convergence to the solution that maximizes social welfare. Finally, we describe several extensions which illustrate the flexibility and tractability of our framework.

1 Introduction

Ad-auctions lie at the heart of search markets and generate billions of dollars in revenue for platforms such as Bing and Google. Sponsored search auctions provide a distributed mechanism where advertisers compete for their adverts to be shown to users of the search platform, by bidding on search terms associated with queries.

The earliest search auction¹ required that advertisers bid a separate price to place an advert in each position on the search page. This design was soon abandoned for one where an advertiser simply bid an amount per click: this amount was converted to adjusted bids for each position by multiplication by the platform's estimate of click-through probabilities; the highest adjusted bid won the first position, the second-highest the second position, and so on, with payments only made when an advert was clicked. The shift from an advertiser making a separate bid for each position to the advertiser making a single bid and being charged per click is an example of *conflation* (Milgrom (2010)): advertisers are required to make the same bid per click whatever the position of the advert.

The design used to assign adverts to positions on the page and the rules used to determine payments have changed several times, with platforms such as BingAds or Google AdWords using variants of the generalised second-price auction (GSP) to determine the price per click: under GSP the amount an advertiser pays when its advert is actually clicked is the smallest price per click that, if bid, would have won the same advert position (Varian (2007); Edelman et al. (2007)).

In current auctions a fundamental information asymmetry between the platform and advertisers has emerged, in that the platform typically knows more than an advertiser about the search being conducted. For example, information on the user conducting the search may comprise location, previous search history, or personal information provided by the user on sign-in to a platform, any of which may affect click-through probabilities. The keyword and additional query information all vary randomly with a distribution that is, in principle, unknown to the platform and advertisers. However, the platform can choose prices and an allocation of adverts to positions using the platform's additional information. In contrast to the platform, the advertiser has to rely on more coarse-grained information, perhaps just the keywords of a query together with a crude categorization of the user. At best an advertiser sees censored information conditional on her advert being shown and clicked: the advertiser has no information about auctions where her advert was either not shown (a losing auction for her) or not clicked, unless the platform chooses to reveal such information.

Variability in the platform's additional information creates additional variability in the observations available to the advertiser. It is difficult for the advertiser to view consecutive allocations by the platform as repeated

¹initiated in 1998 by GoTo.com, which later became Overture and then part of Yahoo!

instances of the same auction since, even if the keyword and range of competitors stay the same, the platform’s additional information varies from search to search. For example, two different searchers for the same keyword may have very different preferences for adverts, giving different click-through probabilities and thus different auctions. These auctions are conflated, with the same bid from an advertiser used in each of them, and this conflation is additional to the conflation over positions. The observations available to the advertiser are inherently stochastic, with the probability of a click-through fluctuating from search to search, and need to be filtered in order to estimate click-through rates. Thus the information asymmetry between platform and advertisers induces a temporal asymmetry: the platform observes each search as it happens, whereas the advertiser has to rely on delayed and aggregated feedback. At the time of writing, platforms typically provide delayed feedback to the advertiser on quantities such as number of impressions (i.e., appearances), average position, average cost per click and so on, averaged over some interval of time.

In this paper we develop a framework to address the information and temporal asymmetries directly. The framework is described in Section 2 and operates as follows. Each advertiser submits a real-valued bid. Using the advertisers’ bids and the platform’s estimated click-through probabilities, the platform assigns adverts to maximize the (expected) bid of an advert receiving a click. This is a classical assignment problem and is solved with low computational overhead on each search instance. We price adverts with a form of parametrized VCG-payment: if an advert receives a click-through, then the pricing of that advert requires one further solution to the assignment problem. We model advertiser i as a utility maximizer, who wants to maximize her payoff $u_i(y_i) = U_i(y_i) - \pi_i y_i$, where U_i is her private, concave, utility function and where π_i and y_i are, respectively, the (average) price she pays per click and the click-through rate she achieves. Under a monotonicity condition which will be satisfied if the platform’s additional information is sufficiently fine-grained, we prove that the Nash equilibrium achieved by advertisers maximizes the social welfare of the advertisers. The framework is extended, in Section 7, to allow an advertiser to make different bids for different keywords or categories of user.

Optimization frameworks of this form are well-established in the communication network community, where users, the network and its components must be separated. There the phrase “Network Utility Maximization” has been coined, but this framework has only recently found its way into Mechanism Design (see Maheswaran and Basar (2004), Yang and Hajek (2007) and Johari and Tsitsiklis (2009)). By contrast, much of the existing literature on sponsored search has needed to restrict attention to an isolated instance of an auction (a single query, repeated without variation) to make progress. We focus upon the stream of search queries: their randomness and the resulting information asymmetry is an intrinsic aspect of our framework.

Implementation of a Nash equilibrium in the economics literature is typically based on the assumption of complete information. In the context of sponsored search, where an advertiser is bidding in a conflated set of auctions with little information on users or competitors (see Pin and Key (2011)), the complete information assumption is not compelling. As Yang and Hajek (2006) discuss in the context of communication networks, an alternative justification for equilibrium is needed and is available. In Section 5 we consider dynamics and convergence under adaptive bid updates by advertisers, and show that under smooth updating of bids, bid trajectories converge to the unique Nash equilibrium.

The mechanism established by our analysis is simple, flexible and implementable. It reduces to the preferred equilibrium of the generalized second price auction (GSP) described in Varian (2007); Edelman et al. (2007) in the special case considered there. But GSP requires an ordered layout of interchangeable adverts, and does not readily adapt to more complex page layouts, such as text rich adverts, adverts of variable size or adverts incorporating images – each of which are of current and increasing demand for modern online advertisement platforms. However, the flexibility to compare and price complex assignments is inherent in VCG mechanisms, and through this, we determine efficient pricing implementations for general page layouts.

It is well known in the economic literature that market clearing prices that equate to marginal utility will maximize social welfare. However, this does not guarantee that such prices can be implemented on the relevant time-scales, where adverts are assigned per impression and charged per click, and search-engine-wide optimization is a highly non-trivial task. Our results show that social welfare can be optimized by a low complexity mechanism which assigns and prices adverts on the time-scales required for sponsored search.

1.1 Outline

In Section 2, we introduce a model of sponsored search where the platform distributes advertisers’ bids over an infinitely large collection of keyword auctions. We define an auction mechanism where an assignment problem is solved for each search occurrence (Section 2.1). We introduce a monotonicity property, requiring click-through rates to continuously increase with bids. The mechanism’s pricing scheme is defined in Section 2.2. We discuss three per-click price implementations, two are deterministic and one is randomized.

In Section 3, we discuss the objective of platform-wide efficiency for a collection of advertisers with concave utility functions. We apply a decomposition argument to a social welfare optimization taken over the uncountably infinite set of constraints in our model of sponsored search. The argument is based on techniques from convex optimization and duality; proofs, complicated by the infinite setting, are mostly relegated to the ap-

pendix. These preliminaries establish that if advertisers equate bids with their marginal utility and the platform solves a maximum weighted matching assignment problem for each search instance then social welfare will be maximized. Crucially, the time-scale and information asymmetry in this decomposition are those relevant to sponsored search. The advertisers are optimizing over a slower time-scale than the platform, and the platform uses the submitted bids to solve on-line a form of generalized first price auction.

After these preliminaries, in Section 4 we make the connection with mechanism design and strategic advertisers. In particular, we find the form of a rebate which incentivizes advertisers to truthfully declare bids that equate to their marginal utilities. This produces a unique Nash equilibrium which implements our decomposition, and this is our main result (Theorem 1). By the separability of our optimization, the rebate can be computed by a simple mechanism requiring a single additional computation, namely the solution of an assignment problem, for each click-through. Hence assignment and pricing occur per search query and involve straightforward polynomial-time computations. The platform solves a (primal) assignment problem on a per-search time-scale; and advertisers maximize their payoffs by solving a dual optimization problem over a longer time-scale (Proposition 4). Finally we prove the theorem, which essentially follows from strong duality.

Section 5 contains our discussion of dynamics and of convergence to the Nash equilibrium under adaptive bid updates by advertisers. In Section 6 we allow more complex page layouts and control of the number of positions displayed (for instance, through reserve prices), and in Section 7 we allow advertisers to make different bids for different keywords or categories of user. Section 8 discusses the relationship between our results and earlier work, and Section 9 concludes.

2 The Assignment and Pricing Model

We begin with notation that reflects a sponsored search setting, where a limited set of adverts are shown in response to users submitting search queries. We let $i \in \mathcal{I}$ index the finite set of advertisers. Each has an advert which they wish to be shown on the pages of search results. An advert, when shown, is placed in a slot $l \in \mathcal{L}$. The set of slots is ordered, with the first (lowest ordered) slot representing the top slot. Let $\tau \in \mathcal{T}$ index the *type* of a search conducted by a user. The set \mathcal{T} is an infinitely large set. The type τ may incorporate information such as the keywords, location, previous search history, and any other information the platform has on the search or searcher. As τ varies, features – such as the keyword – are allowed to change. Let p_{il}^τ be the probability of a click-through on advert i if is shown in slot l : this probability is estimated by the platform and will depend on the type τ .

Over time, a large number of searches from the set \mathcal{T} are made. We assume these occur with distribution \mathbb{P}_τ . Thus we view the click-through probability $p_{il} : \mathcal{T} \rightarrow [0, 1]$ as a random variable defined on the type space \mathcal{T} and with distribution \mathbb{P}_τ . For example, the random variables $p = (p_{il} : i \in \mathcal{I}, l \in \mathcal{L})$ might admit a joint probability density function $f(p)$. So, for $z = (z_{il} : i \in \mathcal{I}, l \in \mathcal{L}) \in [0, 1]^{\mathcal{I} \times \mathcal{L}}$,

$$\mathbb{P}_\tau(p \leq z) = \int_{[0,1]^{\mathcal{I} \times \mathcal{L}}} \mathbf{I}[p \leq z] f(p) dp.$$

Here \mathbf{I} is the indicator function and vector inequalities, e.g., $p \leq z$, are taken componentwise, $p_{il} \leq z_{il} \forall i \in \mathcal{I}, l \in \mathcal{L}$.

We exploit the inherent randomness in p_{il} for the optimal placement of adverts. We assume that the platform has access to the information about the query captured in τ , and so can successfully predict the click-through probability p_{il}^τ , whilst the advertiser does not have access to such fine-grained search information. Later, in Sections 4 and 5, we shall see that the platform can use this information asymmetry to guide the auction system towards an optimal outcome.

2.1 Assignment Model

Next we describe a mechanism by which the platform assigns adverts. Suppose advertiser i submits a bid b_i , which reflects what the advertiser is willing to pay for a click-through. The bid b_i is a non-negative real number. Later, in Section 7, we shall allow an advertiser to submit different bids for different categories of search type, for example for different keywords.

Let $b = (b_i, i \in \mathcal{I})$. Given the information (τ, b) , the following optimization maximizes the expected sum of bids on click-throughs from a single search.

ASSIGNMENT(τ, b)

$$\text{Maximize} \quad \sum_{i \in \mathcal{I}} b_i \sum_{l \in \mathcal{L}} p_{il}^\tau x_{il}^\tau \quad (2.1a)$$

$$\text{subject to} \quad \sum_{i \in \mathcal{I}} x_{il}^\tau \leq 1, \quad l \in \mathcal{L}, \quad (2.1b)$$

$$\sum_{l \in \mathcal{L}} x_{il}^\tau \leq 1, \quad i \in \mathcal{I}, \quad (2.1c)$$

$$\text{over} \quad x_{il}^\tau \geq 0, \quad i \in \mathcal{I}, l \in \mathcal{L}. \quad (2.1d)$$

The above optimization is an assignment problem, where the constraint (2.1b) prevents a slot containing more than one advert, and the constraint (2.1c) prevents any single advert being shown more than once on a search page. The assignment problem is highly appealing from a computational perspective, firstly, because an integral solution can be found efficiently (see Kuhn (1955); Bertsekas (1988)) and, secondly, because there is no need to pre-compute the assignment. The assignment problem can be solved on each occurrence of a search of type $\tau \in \mathcal{T}$, and an integral solution forms a maximum weighted matching of advertisers \mathcal{I} with slots \mathcal{L} .

We apply the convention that if $b_i = 0$ then $x_{il}^\tau = 0$ for $l \in \mathcal{L}$, so that a zero bid does not receive clicks. Let

$$y_i^\tau = \sum_{l \in \mathcal{L}} p_{il}^\tau x_{il}^\tau, \quad y_i = \mathbb{E}_\tau y_i^\tau. \quad (2.2)$$

The solution x^τ to the assignment problem (2.1) may not be unique: however the solution will be unique with probability one if, for example, the distribution of click-through probabilities p admits a density. We make the milder assumption that $(y_i^\tau, i \in \mathcal{I})$ is unique with probability one.

Note that y_i^τ is the click-through rate for advertiser i from a given search page, and y_i is the click-through rate averaged over \mathcal{T} . (We shall *not* use y_i for the random variable y_i^τ .) In our model the information asymmetry between the platform and advertiser is captured by the search type τ which is known to the platform but not to the advertiser: thus we assume that y_i^τ is known to the platform, from its solution to the assignment problem, while only the average y_i is reported to, or accessible for estimation by, advertiser i . For an optimal solution to the above assignment problem, write $y_i^\tau = y_i^\tau(b)$ to emphasize the dependence of y_i^τ on the vector of bids b and, similarly, write $y_i = y_i(b)$. Let (b'_i, b_{-i}) be the vector obtained from b by replacing the i th component by b'_i .

We shall assume the following *monotonicity property* of solutions of ASSIGNMENT(τ, b). We assume that $y_i(b_i, b_{-i})$ takes the value 0 when $b_i = 0$, and is strictly increasing in b_i and continuous in (b_i, b_{-i}) whenever any component of b_{-i} is positive. Without the monotonicity property $y_i(b)$ will be increasing in b_i but may not be strictly increasing or continuous. A similar assumption has been made by Nekipelov et al. (2015), who argue that the assumption is natural and satisfied by the sponsored search data they analyze.

The monotonicity property will generally follow from sufficient variability of click-through rates. For instance, a sufficient condition is that the random variables p admit a continuous density $f(p)$ on the set of click-through probabilities $\tilde{\mathcal{P}} = \{p \in [0, 1]^{|\mathcal{I}| \times |\mathcal{L}|} : p_{il} \geq p_{ik}, l < k\}$ which is positive on a neighborhood containing the origin. Observe that on $\tilde{\mathcal{P}}$ the click-through probability for a given advert increases as the slot it is shown in decreases. The following result establishes the monotonicity property under the above sufficient condition.

Proposition 1. *If the distribution \mathbb{P}_τ admits a continuous probability density function on $\tilde{\mathcal{P}}$ which is positive on a neighborhood containing the origin then the mapping $b_i \mapsto y_i(b_i, b_{-i})$ satisfies the monotonicity property.*

The proof of Proposition 1 is given in Appendix A. The sufficient condition of Proposition 1 is far from necessary, as we shall illustrate later in Example 1. Our earlier assumption that $(y_i^\tau, i \in \mathcal{I})$ is unique with probability one is implied by the monotonicity property.

2.2 Pricing Model

Once adverts are allocated, prices must be determined for any resulting click-throughs. We consider a mechanism where the expected rate of payment by advertiser i is

$$\pi_i(b) y_i(b) = \int_0^{b_i} \left(y_i(b) - y_i(b'_i, b_{-i}) \right) db'_i. \quad (2.3)$$

Here, as before, $b = (b_i : i \in \mathcal{I})$ is the vector of advertisers' bids and $y_i(b)$ is the resulting click-through rate for advertiser i . We discuss later the rationale for this formula in the proper context of mechanism design, in Section 4. For now we note that the rate of payment (2.3) can be readily implemented by the platform at a low computational cost. We give three examples of implementations: the first uses randomization to estimate the integral (2.3); the second is a form of VCG price; and the third uses the solution of a linear program due to Leonard (1983). The first two require the solution of just one additional instance of the assignment problem per click-through.

A randomized price.

Suppose the platform solves $\text{ASSIGNMENT}(\tau, b)$, and observes a click-through on (i, l) — that is the solution has $x_{il}^\tau = 1$, and the user clicks on the advert in position l , which is for advertiser i . To price this advert, the platform chooses b'_i uniformly and randomly on the interval $(0, b_i)$ and additionally solves $\text{ASSIGNMENT}(\tau, (b'_i, b_{-i}))$. Let $y_i^\tau(b'_i, b_{-i}) = \sum_{l \in \mathcal{L}} p_{il}^\tau x_{il}^\tau$ under a solution to this problem. The platform then charges advertiser i an amount

$$b_i \left(1 - \frac{y_i^\tau(b'_i, b_{-i})}{y_i^\tau(b)} \right) \quad (2.4)$$

for the click-through. This charge does not depend on the distribution \mathbb{P}_τ , and will lie between 0 and b_i . Taking expectations over τ and b'_i shows that the expected rate of payment by advertiser i is

$$\mathbb{E}_{\tau, b'_i} \left[\sum_{l \in \mathcal{L}} p_{il}^\tau x_{il}^\tau b_i \left(1 - \frac{y_i^\tau(b'_i, b_{-i})}{y_i^\tau(b)} \right) \right] = b_i (y_i(b) - \mathbb{E}_{b'_i} [y_i(b'_i, b_{-i})]) = \int_0^{b_i} (y_i(b) - y_i(b'_i, b_{-i})) db'_i,$$

recovering expression (2.3).

Observe that the additional instance of the assignment problem does not determine the assignment, and thus will not slow down the page impression: rather, it is used to calculate the charge (2.4) for a click-through. Indeed, one could imagine a charge b_i on the click-through, followed by a later rebate of a proportion $y_i^\tau(b'_i, b_{-i})/y_i^\tau(b)$ of the charge. The rebate depends on the uniform random variable b'_i as well as the random variable τ : next we shall see that we can remove the dependence on b'_i .

A parametrized VCG price.

Note that

$$\int_0^{b_i} y_i(b'_i, b_{-i}) db'_i = \sum_j b_j y_j(b) - \sum_{j \neq i} b_j y_j(0, b_{-i}),$$

since both expressions share the same derivative with respect to b_i (see Proposition 2 of Appendix A) and both expressions take the value 0 when $b_i = 0$. Thus the rate of payment (2.3) can be implemented by a charge b_i on a click-through followed by a later rebate

$$\frac{1}{y_i^\tau(b)} \left(\sum_j b_j y_j^\tau(b) - \sum_{j \neq i} b_j y_j^\tau(0, b_{-i}) \right). \quad (2.5)$$

The rebate calculation again requires the solution of one additional instance of the assignment problem, this time omitting advertiser i . This calculation is familiar as the VCG mechanism when the utility function for advertiser j , $j \in \mathcal{I}$, is replaced by the surrogate linear utility $b_j y_j$. The charge minus the rebate has the usual VCG interpretation as the externality caused by advertiser i , but under these surrogate utilities.

Computing all prices simultaneously.

Leonard (1983) has shown that VCG prices in assignment games are a minimal solution to a dual assignment problem, and this allows prices for all potential click-throughs to be calculated from the solution to just one optimization problem.

Let A^τ be the maximal value achieved by the objective function (2.1a) in the assignment problem (2.1). Then per-impression VCG prices are given by the solution v^τ, s^τ to the following optimization problem.

$$\begin{array}{ll} \text{Minimize} & \sum_{l \in \mathcal{L}} v_l \\ \text{subject to} & \sum_{i \in \mathcal{I}} s_i + \sum_{l \in \mathcal{L}} v_l = A^\tau, \\ & s_i + v_l \geq b_i p_{il}^\tau, \quad i \in \mathcal{I}, l \in \mathcal{L}, \\ \text{over} & s_i \geq 0, v_l \geq 0, \quad i \in \mathcal{I}, l \in \mathcal{L}. \end{array}$$

An initial feasible solution to this dual assignment program is given by the dual variables corresponding to an optimum of the assignment problem (2.1) and techniques for its solution are reviewed in Bikhchandani et al. (2002).

This formulation allows for either pay-per-click or pay-per-impression pricing of adverts. After solving the problem for v^τ, s^τ , advertiser i can either be charged the price v_l^τ for an impression of her advert in slot l or be charged the price $v_l^\tau/p_{il}^\tau = b_i - s_i^\tau/p_{il}^\tau$ on a click-through: in the latter case the result of Leonard (1983) implies

that the rebate s_i^τ/p_{i1}^τ will equal expression (2.5). Observe that the dual assignment problem to be solved is identical whichever advert is clicked on.

We end this section with a setting where particularly simple closed forms are available for prices.

Example 1. *If there is a single slot then the slot will be assigned to the bidder i with the highest value of $b_i p_{i1}^\tau$, and if this results in a click-through then the charge will be $\max_{j \neq i} b_j p_{j1}^\tau / p_{i1}^\tau$, a second price auction on the products $b_j p_{j1}^\tau$.*

Suppose next there are L slots with I advertisers bidding and further suppose that the click-through probabilities take the form $p_{il}^\tau = q_i^\tau p_l$ where $p_1 > p_2 > \dots > p_L$. Here p_l is a slot effect, and q_i^τ is an advertiser effect which may depend on the search query (for example, it may depend on some measure of distance between the searcher and the advertiser). Define the search-adjusted bid $b_i^\tau = b_i q_i^\tau$ and, given τ , order the advertisers so that $b_1^\tau > b_2^\tau > \dots > b_I^\tau$. Then advertisers $1, 2, \dots, \min\{L, I\}$ are allocated slots $1, 2, \dots, \min\{L, I\}$ respectively. If necessary ties can be broken randomly.

In this example it is straightforward to calculate the expected value of expression (2.4) over b_i^τ explicitly. Set $p_{L+1} = 0$ and $b_i = b_i^\tau = 0$ for $i > I$. Upon a click-through on slot l advertiser l is charged the amount π_l^τ where

$$\pi_l^\tau q_l^\tau = b_{l+1}^\tau - \frac{1}{p_l} \sum_{m=l+1}^L p_m (b_m^\tau - b_{m+1}^\tau), \quad l = 1, 2, \dots, L.$$

Expressed as a recursion this implies

$$\pi_l^\tau = \frac{q_{l+1}^\tau}{q_l^\tau} \left(b_{l+1} - \frac{p_{l+1}}{p_l} (b_{l+1} - \pi_{l+1}^\tau) \right), \quad l = 1, 2, \dots, L \quad (2.6)$$

recovering an equilibrium of the generalized second price auction, Edelman et al. (2007). Note, however, that the charges (2.6), and indeed the slots allocated, fluctuate with the search type τ . The expected revenue, given τ , is

$$\sum_{m=1}^L \pi_m^\tau q_m^\tau p_m = \sum_{m=1}^L (p_m - p_{m+1}) b_{m+1} q_{m+1}^\tau. \quad (2.7)$$

In the model considered by Edelman et al. (2007) and Varian (2007) the random variables $(q_i^\tau, i \in \mathcal{I})$ are all in fact constants, and in this case there may be multiple Nash equilibria. For example, suppose $L = I = 2$: then for either one of the advertisers to bid very high and the other to bid very low is a Nash equilibrium. We shall see in following sections that provided our monotonicity condition is satisfied there is a unique Nash equilibrium.

The restriction that click-through probabilities have the product-form $p_{il}^\tau = q_i^\tau p_l$ implies they lie in a linear subspace of $\tilde{\mathcal{P}}$: thus they do not have a density over $\tilde{\mathcal{P}}$, and so we cannot appeal to Proposition 1 to justify the monotonicity property, in particular that $y_i(b)$ is a strictly increasing and continuous function of b_i . But if the advertiser effects $(q_i^\tau, i \in \mathcal{I})$ have a continuous probability density positive on a neighbourhood of the origin in $\{q \in [0, 1]^{|\mathcal{I}|\}$ then the monotonicity property will follow. Essentially the variability of the advertiser effect q_i^τ smooths out the impact of the bid b_i sufficiently that the rate $y_i(b_i, b_{-i})$ is continuous in b_i .

3 Optimization Preliminaries

In this section we present an optimization problem which we use to develop various decomposition and duality results. In particular, we find that if advertisers equate bids with their marginal utility and the platform solves a maximum weighted matching assignment problem for each search instance, then social welfare will be maximized.

We suppose each advertiser i has a utility function, $y_i \mapsto U_i(y_i)$, where $U_i(\cdot)$ is non-negative, increasing, strictly concave and continuously differentiable. Our objective is to place adverts so as to maximize the sum of these utilities, in other words to maximize social welfare. To simplify the statement of results we shall assume further that $U_i'(y_i) \rightarrow \infty$ as $y_i \downarrow 0$ and $U_i'(y_i) \rightarrow 0$ as $y_i \uparrow \infty$. The maximization of social welfare by the auction system is the following problem.

SYSTEM($U, \mathcal{I}, \mathbb{P}_\tau$)

$$\text{Maximize} \quad \sum_{i \in \mathcal{I}} U_i(y_i) \quad (3.1a)$$

$$\text{subject to} \quad y_i = \mathbb{E}_\tau \left[\sum_{l \in \mathcal{L}} p_{il}^\tau x_{il}^\tau \right], \quad i \in \mathcal{I}, \quad (3.1b)$$

$$\sum_{i \in \mathcal{I}} x_{il}^\tau \leq 1, \quad l \in \mathcal{L}, \tau \in \mathcal{T}, \quad (3.1c)$$

$$\sum_{l \in \mathcal{L}} x_{il}^\tau \leq 1, \quad i \in \mathcal{I}, \tau \in \mathcal{T}, \quad (3.1d)$$

$$\text{over} \quad x_{il}^\tau \geq 0, y_i \geq 0 \quad i \in \mathcal{I}, l \in \mathcal{L}. \quad (3.1e)$$

Inequalities (3.1c) and (3.1d) are just the scheduling constraints (2.1b) and (2.1c), that each slot can show at most one advert and that each advertiser can show at most one advert, while equality (3.1b) recaps the definition (2.2) of y_i , the expected click-through rate. Over these constraints we maximize social welfare, i.e., the aggregate sum of the utilities.

To solve the above optimization, one could imagine that there is a centralized designer who knows everything about the entire system: the advertisers' utilities $U_i(\cdot), i \in \mathcal{I}$, click-through probabilities $p_{il}^\tau, i \in \mathcal{I}, l \in \mathcal{L}, \tau \in \mathcal{T}$, and the distribution \mathbb{P}_τ over these probabilities. This designer then attempts to assign adverts in a way so that $y_i, i \in \mathcal{I}$, the click-through rates received by advertisers, maximize social welfare. The solution of such an optimization by centralized means is not possible — for example, the utilities will not be known — but the form of the solution will help us develop an appropriate decomposition, respecting the time-scales relevant to the platform and advertisers. In the next section, on mechanism design, we consider the game theoretic aspects that arise when, instead of a single system optimizer, the platform and advertisers have differing information and incentives.

Incorporating the constraint (3.1b) into the objective function (3.1a) gives the Lagrangian

$$L_{sys}(x, y; b) = \sum_{i \in \mathcal{I}} U_i(y_i) + \sum_{i \in \mathcal{I}} b_i \mathbb{E}_\tau \left[\sum_{l \in \mathcal{L}} p_{il}^\tau x_{il}^\tau - y_i \right],$$

where $b_i, i \in \mathcal{I}$ are the Lagrange multipliers associated with the constraints (3.1b), with $b_i \geq 0$. Notice, we intentionally omit the scheduling constraints from our Lagrangian. Thus we seek to maximize the Lagrangian subject to the constraints (3.1c-3.1d) as well as (3.1e). Let \mathcal{S} be the set of variables $x^\tau = (x_{il}^\tau : i \in \mathcal{I}, l \in \mathcal{L})$ satisfying the assignment constraints (2.1b-2.1d), and let \mathcal{A} be the set of variables $x = (x^\tau \in \mathcal{S} : \tau \in \mathcal{T})$ satisfying the assignment constraints (3.1c-3.1e). We see that our Lagrangian problem is separable in the following sense

$$\max_{x \in \mathcal{A}, y \geq 0} L_{sys}(x, y; b) = \sum_{i \in \mathcal{I}} \max_{y_i \geq 0} \{U_i(y_i) - b_i y_i\} \quad (3.2a)$$

$$+ \mathbb{E}_\tau \left[\max_{x^\tau \in \mathcal{S}} \sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{L}} b_i p_{il}^\tau x_{il}^\tau \right]. \quad (3.2b)$$

Define

$$U_i^*(b_i) = \max_{y_i \geq 0} \{U_i(y_i) - b_i y_i\}. \quad (3.3)$$

The optimization over y_i contained in the definition (3.3) would arise if advertiser i were presented with a fixed price per click-through of b_i : if allowed to choose freely her click-through rate, she would then choose y_i such that $U_i'(y_i) = b_i$. By our assumptions on $U_i(\cdot)$, this equation has a unique solution for all $b_i \in (0, \infty)$. Call $D_i(\xi) = \{U_i'\}^{-1}(\xi)$ the *demand* of advertiser i at price ξ . It follows that $U_i^*(b_i)$ can be written in the form

$$U_i^*(b_i) = \int_{b_i}^{\infty} D_i(\xi) d\xi; \quad (3.4)$$

call this advertiser i 's *consumer surplus* at the price b_i . From this expression we can deduce that $U_i^*(b_i)$ is positive, decreasing, strictly convex and continuously differentiable.

Observe that the maximization inside the expectation (3.2b) is simply the problem ASSIGNMENT(τ, b), and thus we can write

$$\max_{x \in \mathcal{A}, y \geq 0} L_{sys}(x, y; b) = \sum_{i \in \mathcal{I}} U_i^*(b_i) + \sum_{i \in \mathcal{I}} b_i y_i(b).$$

The Lagrangian dual of the SYSTEM problem (3.1) can thus be written as follows.

DUAL(U^*, y, \mathcal{I})

$$\text{Minimize} \quad \sum_{i \in \mathcal{I}} (U_i^*(b_i) + b_i y_i(b)) \quad (3.5a)$$

$$\text{over} \quad b_i \geq 0, \quad i \in \mathcal{I}. \quad (3.5b)$$

Owing to the size of the type space \mathcal{T} , the optimization (3.1) has a potentially uncountable number of constraints. This presents certain technical difficulties, for instance those associated with proving strong duality. These issues are dealt with in the appendix, where the proofs of the following two propositions are presented.

We first observe that the SYSTEM problem decomposes into optimizations relevant to the advertisers and to the platform.

Proposition 2 (Decomposition). *Variables \tilde{y} , \tilde{x}^τ , $\tau \in \mathcal{T}$, satisfying the feasibility conditions (3.1b-3.1e) are optimal for $SYSTEM(U, \mathcal{I}, \mathbb{P}_\tau)$ if and only if there exist \tilde{b}_i , $i \in \mathcal{I}$, such that*

A. \tilde{b}_i minimizes $U_i^*(b_i) + b_i \tilde{y}_i$ over $b_i \geq 0$, for each $i \in \mathcal{I}$,

B. \tilde{x}^τ solves $ASSIGNMENT(\tau, \tilde{b})$ with probability one under the distribution \mathbb{P}_τ over $\tau \in \mathcal{T}$.

In this proposition, the optimization in Condition A does not naturally correspond to the bidding behavior of strategic advertisers, at least in its present form. Hence we need to examine the implications of Condition A for the construction of prices (2.3) that do give strategic advertisers the incentive to solve the SYSTEM problem. We do this in the next section, Section 4. There we shall also see that the per-click pricing implementations (2.4) and (2.5) are made possible by the decomposition into per-impression assignments, Condition B.

The optimal bids \tilde{b} can be further understood through the following dual characterization.

Proposition 3 (Dual Optimality).

a) *The objective of the dual problem (3.5) is continuously differentiable for $b > 0$ and is minimized uniquely by the positive vector $\tilde{b} = (\tilde{b}_i : i \in \mathcal{I})$ satisfying, for each $i \in \mathcal{I}$,*

$$\frac{dU_i^*}{db_i}(\tilde{b}_i) + y_i(\tilde{b}) = 0. \quad (3.6)$$

b) *If \tilde{b} is an optimal solution to the DUAL problem (3.5) then $x^\tau(\tilde{b})$, $y(\tilde{b})$ are optimal for the SYSTEM problem (3.1).*

The dual provides a finite parameter optimization from which the SYSTEM problem can be solved. Moreover, (3.6) provides conditions on advertiser demands which, to solve the SYSTEM problem, must be effected by the auction system in strategic form.

4 Mechanism Design

We now prove that our mechanism implements our system optimization. In the last section we demonstrated how this global problem can be decomposed into two types of sub-problem: one, where the platform finds an optimal assignment given click-through probabilities; and the other, where the dual variables b are each set to solve a certain single parameter dual problem. In this section we suppose the advertisers act strategically, anticipating the result of the platform's assignment and attempting to maximize their payoff.

Henceforth b_i is the *bid* submitted by advertiser i and, as a function of these bids, we formulate prices that incentivize the advertisers to choose bids that result in an assignment that solves the SYSTEM problem (3.1).

Consider a mechanism where, given the vector of bids $b = (b_i : i \in \mathcal{I})$, each advertiser, i , receives a click-through rate $y_i(b)$, and from this derives a benefit $U_i(y_i(b))$ and is charged an expected price $\pi_i(b)$ per click. The payoff to advertiser i arising from a vector of bids $b = (b_i : i \in \mathcal{I})$ is then

$$u_i(b) = U_i(y_i(b)) - \pi_i(b)y_i(b). \quad (4.1)$$

A *Nash equilibrium* is a vector of bids $b^* = (b_i^* : i \in \mathcal{I})$ such that, for $i \in \mathcal{I}$ and all b_i

$$u_i(b^*) \geq u_i(b_i, b_{-i}^*). \quad (4.2)$$

Here (b_i, b_{-i}^*) is obtained from the vector b^* by replacing the i th component by b_i .

The main result of this section is the following.

Theorem 1. *If prices are charged so that the expected rate of payment by advertiser i , for $i \in \mathcal{I}$, is given by expression (2.3) then there exists a unique Nash equilibrium, and it is given by the vector of optimal prices identified in Proposition 3. Thus the assignments achieved at the Nash equilibrium, $x^\tau(b^*), y(b^*)$, form a solution to the SYSTEM problem (3.1).*

The result states that, given adverts are assigned according to the assignment problem (2.1), the game theoretic equilibrium reached by advertisers attempting to maximize their respective payoffs u_i solves the problem $\text{SYSTEM}(U, \mathcal{I}, \mathbb{P}_\tau)$. Since $y_i(b'_i, b_{-i})$ is a strictly increasing function of the bid b'_i , it follows from (2.3) that the price $\pi_i(b)$ must be strictly lower than the bid b_i . Setting a price lower than the submitted bid is a prevalent feature of online auctions used by search engines, and, as we emphasized in Section 2, the prices (2.3) can be practically implemented in a sponsored search setting.

We note that, in this section, each advertiser expresses their preferences through a single bid. This framework extends naturally to the case where advertisers place multiple bids over multiple different keywords (or search categories). This extension is given in Section 7.

4.1 Proof of Theorem 1

To establish Theorem 1 we will require an additional result, Proposition 4, which indicates how maximal payoffs achieved by each advertiser relate to the solution of the dual problem, given by Proposition 3 from the previous section.

Proposition 4 (Mechanism Dual). *For each positive choice of $b_{-i} = (b_j : j \neq i, j \in \mathcal{I})$, the following equality holds*

$$\max_{b_i \geq 0} u_i(b) = \min_{b_i \geq 0} \left\{ U_i^*(b_i) + \int_0^{b_i} y_i(b'_i, b_{-i}) db'_i \right\}. \quad (4.3)$$

Moreover, the optimizing b_i for both expressions is the same, is unique and finite, and satisfies

$$\frac{d}{db_i} U_i^*(b_i) + y_i(b) = 0. \quad (4.4)$$

Proof. We calculate the conjugate dual of the payoff function (4.1). Let $P_i(y_i)$ be the function whose Legendre-Fenchel transform is

$$P_i^*(b_i) = \int_0^{b_i} y_i(b'_i, b_{-i}) db'_i.$$

The above function is increasing and convex, and we know from Fenchel's Duality Theorem (Borwein and Lewis, 2006, Theorem 3.3.5) that

$$\max_{y_i \geq 0} \{U_i(y_i) - P_i(y_i)\} = \min_{b_i \geq 0} \{U_i^*(b_i) + P_i^*(b_i)\}. \quad (4.5)$$

Next we calculate the function P_i from the dual of the function P_i^* above. By the Fenchel-Moreau Theorem, Borwein and Lewis (2006), we know this to be

$$P_i(y_i) = \min_{b_i \geq 0} \left\{ b_i y_i - \int_0^{b_i} y_i(b'_i, b_{-i}) db'_i \right\}.$$

The optimum in this expression occurs when $y_i(b) = y_i$. Substituting this back, since $b_i \mapsto y_i(b)$ is strictly increasing, we have that

$$P_i(y_i) = \int_0^\infty (y_i - y_i(b'_i, b_{-i})) \mathbf{I}[y_i(b'_i, b_{-i}) \leq y_i] db'_i. \quad (4.6)$$

In other words, as expected with the Legendre-Fenchel transform, the area under the curve $y_i(b_i, b_{-i})$ is converted to the area to the left of the curve $y_i(b_i, b_{-i})$. Further, notice, if $y_i > \max_{b_i} y_i(b_i, b_{-i})$ then $P_i(y_i) = \infty$, and thus the finite range of the function $y_i \mapsto P_i(y_i)$ is exactly the same as that of $b_i \mapsto P_i(y_i(b))$. Noting (4.6) and this last observation, the equality (4.5) now reads

$$\begin{aligned} \min_{b_i \geq 0} \left\{ U_i^*(b_i) + \int_0^{b_i} y_i(b'_i, b_{-i}) db'_i \right\} &= \max_{y_i \geq 0} \{U_i(y_i) - P_i(y_i)\} \\ &= \max_{b_i \geq 0} \{U_i(y_i(b)) - P_i(y_i(b))\} \\ &= \max_{b_i \geq 0} \{U_i(y_i(b)) - \pi_i(b) y_i(b)\}. \end{aligned}$$

In the final equality we note from the definition (2.3) that $P_i(y_i(b)) = \pi_i(b)y_i(b)$. This gives the equality (4.3).

We now show that both expressions (4.3) are determined at the same unique value of b_i . The function $U_i^*(b_i) - P_i^*(b)$ is a strictly convex differentiable function of b_i , whose unique minimum is given by the required expression (4.4). Further, $b_i \mapsto y_i(b_i, b_{-i})$ is strictly increasing and b_i achieves the range of the strictly concave function $U_i(y_i) - P_i(y_i)$ under $y_i = y_i(b_i, b_{-i})$. Thus $U_i(y_i) - P_i(y_i)$ is maximized uniquely by $y_i = y_i(b_i, b_{-i})$ (and thus uniquely by b_i) satisfying

$$\frac{d}{dy_i}U_i(y_i(b)) - b_i = 0. \quad (4.7)$$

Since $\frac{d}{db_i}U_i^*$ is the inverse of the strictly increasing function $\frac{d}{dy_i}U_i$, it is clear that (4.4) and (4.7) are equivalent and satisfied by the same unique b_i . This completes the proof. \square

The proof of Theorem 1 follows by observing the optimality conditions of Propositions 3 and 4.

Proof of Theorem 1. Before proceeding with the main argument, we note that a Nash equilibrium must be achieved by positive values of b_i . By applying the mean value theorem, for some \tilde{y} satisfying $0 = y_i(0, b_{-i}) \leq \tilde{y} \leq y_i(b_i, b_{-i})$, we have

$$\begin{aligned} u_i(b_i, b_{-i}) &\geq U_i(0) + U_i'(\tilde{y})(y_i(b_i, b_{-i}) - y_i(0, b_{-i})) - \int_0^{b_i} (y_i(b_i, b_{-i}) - y_i(0, b_{-i}))db'_i \\ &= u_i(0, b_{-i}) + (U_i'(\tilde{y}) - b_i)(y_i(b_i, b_{-i}) - y_i(0, b_{-i})) \\ &> u_i(0, b_{-i}). \end{aligned} \quad (4.8)$$

The second term in (4.8) is positive for b_i sufficiently small, since $U_i'(\tilde{y}) - b_i \nearrow \infty$ as $b_i \searrow 0$ and from our monotonicity property $y_i(b_i, b_{-i}) > y_i(0, b_{-i})$. From this we see that a Nash equilibrium can only be achieved with $b_i > 0$ for each $i \in \mathcal{I}$.

By Proposition 4, $b = (b_i : i \in \mathcal{I}) > 0$ is a Nash equilibrium if and only if condition (4.4) is satisfied for each $i \in \mathcal{I}$. But by Proposition 3b), these conditions hold if and only if b is the unique solution to the dual to the SYSTEM problem. So, the set of Nash equilibria are the optimal prices defined for the decomposition, Proposition 2. By Proposition 3b), the assignment achieved by Nash equilibrium bids maximizes the utilitarian objective $\text{SYSTEM}(U, \mathcal{I}, \mathbb{P}_\tau)$. Finally, by Strong Duality (Theorem 3 of Appendix C), there exists b^* which optimizes the dual problem (3.5), and thus there must be a Nash equilibrium. \square

Remark 1. *The optimality condition (3.6) or (4.4) states that each advertiser's demand, $D_i(b_i)$, and supply, $y_i(b)$, should equate, and is a consequence of the Envelope Theorem. A more familiar context for this form of result is Vickrey pricing (Vickrey (1961)) and Myerson's Lemma (or the Revenue Equivalence Theorem), see Myerson (1981) and (Milgrom, 2004, Theorem 3.3), which are also consequences of the Envelope Theorem. But observe that we are using general utilities, which despite the single input parameter b_i , takes us out of a single parameter type space to which Myerson's Lemma generally applies.*

We have assumed throughout the monotonicity property, ensuring that the mapping $b_i \mapsto y_i(b_i, b_{-i})$ is strictly increasing and continuous. A natural question concerns whether the monotonicity property can be relaxed.

Example 2. *If the mapping is discontinuous, there may be inefficient Nash equilibria, and the $L = I = 2$ case discussed in Example 1, with two advertisers and two slots, provides an illustration. The same difficulty can arise even if the mapping is continuous but not strictly increasing, as we now show. Amend the illustration, by supposing that the advertiser effects q_1^τ, q_2^τ are independent random variables with continuous probability density functions each supported on the interval $(q - \epsilon, q + \epsilon)$ for $q \gg \epsilon > 0$. The mapping $b_i \mapsto y_i(b_i, b_{-i})$ is now continuous, although not strictly increasing. The inefficient Nash equilibria remain, where one of the advertisers bids very high and the other very low. If we assume the densities of q_1^τ, q_2^τ are positive in a neighbourhood of the origin, then the mapping is necessarily strictly increasing, because a small increase in an advertiser's bid will have a small but positive probability of improving the slot allocated to the advertiser: competition exists between the advertisers, whatever their bids, for at least some search types τ , and this ensures the uniqueness and efficiency of the Nash equilibrium.*

5 Dynamics and Convergence

We have seen in Section 2 that our assignment model involves the rapid solution of a computationally straightforward problem for each individual search. The challenge facing an advertiser is of a different form: she has to rely on noisy and possibly delayed feedback averaged over some period of time in order to learn the mean click-through rate y_i that has been achieved by her bid b_i , and she then has to decide whether to vary her bid.

We shall formulate the advertiser's problem in continuous time, and the natural question is whether multiple advertisers smoothly varying their bids $b_i(t)$ as a consequence of their current click-through rates $y_i(t)$ will converge to the Nash equilibrium.

Convergence may not be possible when the search space is discrete, e.g., for an auction on a single search type. Essentially, the search engine does not have enough additional information from the search type τ to fine tune its discrimination between advertisers. However, in sponsored search, there is inherent variability in the search type τ which will influence the click-through probabilities of the advertiser. This is the motivation for our assumption of the monotonicity property, that the distribution \mathbb{P}_τ over \mathcal{T} is such that the click-through rate $y_i(b)$ is a continuous, strictly increasing function of b_i . We shall see that, under models of advertiser response, we are then able to deduce convergence towards a system optimum.

Recall the objective function for the dual of the system problem as derived in Proposition 3,

$$\mathcal{V}(b) = \sum_{i \in \mathcal{I}} U_i^*(b_i) + \sum_{i \in \mathcal{I}} b_i y_i(b). \quad (5.1)$$

This expression is the sum of the consumer surpluses and the revenue achieved by the platform at prices b and, when b is optimal, it is equal to the maximal total welfare as defined by the SYSTEM problem (3.1). Further, $\mathcal{V}(b)$ is continuously differentiable for $b > 0$ with

$$\frac{\partial \mathcal{V}}{\partial b_i} = -D_i(b_i) + y_i(b).$$

We next model advertisers' responses to their observation of click-through rates. We suppose advertiser i changes her bid $b_i(t)$ smoothly (i.e., continuously and differentiably) as a consequence of her observation of her current click-through rate $y_i(t)$ so that

$$\frac{d}{dt} b_i(t) \geq 0 \text{ according as } b_i(t) \leq U_i'(y_i(b(t))). \quad (5.2)$$

This is a natural dynamical system representation of advertiser i varying b_i smoothly in order to improve her payoff $u_i(b)$, given by expression (4.1), under prices (2.3), since under the monotonicity condition a small positive change in b_i will cause a small positive change in $y_i(b)$ and the impact on $u_i(b)$ will be positive or negative as in relation (5.2) - see Lemma ?? in Appendix ?. Note that from the definition of the demand function $D_i(\cdot)$,

$$y_i \leq D_i(b_i) \text{ according as } b_i \leq U_i'(y_i). \quad (5.3)$$

The payoff $u_i(b)$ is maximized over b_i when b_i and $U_i'(y_i(b))$ equate, or equivalently, when $y_i(b)$ and $D_i(b_i)$ equate.

Theorem 2 (Convergence of Dynamics). *Starting from any point $b(0)$ in the interior of the positive orthant, the trajectory $(b(t) : t \geq 0)$ of the above dynamical system converges to a solution of the DUAL problem (3.5). Thus $y(b(t))$, the assignment achieved by the prices $b(t)$, converges to a solution of the SYSTEM problem (3.1).*

Proof. We prove that the objective of the dual problem $\mathcal{V}(b)$, defined above, is a Lyapunov function for the dynamical system. Note that $\mathcal{V}(b)$ is continuously differentiable for $b > 0$. Since $y_i(b) \downarrow 0$ as $b_i \downarrow 0$ and $U'(0) > 0$ it follows from (5.2) that there exists $\delta > 0$, possibly depending on $b_{-i}(t)$, such that $\frac{d}{dt} b_i(t) > 0$ if $b_i(t) \leq \delta$. We deduce that the paths of our dynamical system $(b(t) : t \geq 0)$ are strictly positive and $\mathcal{V}(b(t))$ is continuously differentiable along these paths. Further, the level sets $\{b : \mathcal{V}(b) \leq \kappa\}$ are compact: this is an immediate consequence of the facts that the functions $U_i^*(b_i)$ are positive and decreasing, and, as proven in Lemma 2, that

$$\lim_{\|b\| \rightarrow \infty} \sum_{i \in \mathcal{I}} b_i y_i(b) = \infty.$$

Differentiating $\mathcal{V}(b(t))$ yields

$$\frac{d}{dt} \mathcal{V}(b(t)) = \sum_{i \in \mathcal{I}} \frac{\partial \mathcal{V}}{\partial b_i} \frac{d}{dt} b_i(t) = - \sum_{i \in \mathcal{I}} (D_i(b_i(t)) - y_i(b(t))) \frac{d}{dt} b_i(t) \leq 0,$$

where the inequality follows from relations (5.2) and (5.3), and is strict unless $D_i(b_i(t)) = y_i(b(t))$ for $i \in \mathcal{I}$. By Lyapunov's Stability Theorem, see (Khalil, 2002, Theorem 4.1), the process $(b(t) : t \geq 0)$ converges to the set of points b^* satisfying, for $i \in \mathcal{I}$, $D_i(b_i) = y_i(b)$. Recall that $\frac{dU_i^*}{db_i} = -D_i(b_i)$ and thus, by Proposition 3(a), the price process $b(t)$ converges to an optimal solution to the dual problem (3.5). By the monotonicity property $y(b)$ is continuous, and thus by Proposition 3(b) the click-through rates $y(b(t))$ converge to an optimal solution for the system problem. \square

In the above discussion we model advertisers that smoothly change their bids over time. However, we remark that other convergence mechanisms could be considered. For instance, since our dual optimization problem is convex and continuously differentiable, we can minimize the dual through a coordinate descent algorithm, where each component b_i is sequentially minimized. Such an algorithm could correspond to a game played sequentially with advertisers iteratively maximizing over b_i their payoff $u_i(b_i, b_{-i})$. Previous work on global convergence to a Nash equilibrium using an assumption of local rather than complete information is described by Yang and Hajek (2006).

The dynamical system of this section allows advertisers' bids to adapt to a non-stationary environment, for example if the set of participating advertisers changes. Note that we have left unexplored the statistical aspects of estimating the click-through rates y_i , although some insights are available from the network utility maximization framework for communication networks (Kelly (2003)). In particular, if the period of time over which the click-through rates y_i are estimated is longer then this will improve the statistical accuracy of the estimation, but will also slow down the rate of adaptation to a changing environment; and even in a stationary environment there is necessarily a trade-off between the speed of convergence to, and the stochastic variability around, the system optimum.

Remark 2. *We have assumed $U_i(\cdot)$ is non-negative, increasing and strictly concave, and is continuously differentiable with boundary conditions $U_i'(y_i) \rightarrow \infty$ as $y_i \downarrow 0$ and $U_i'(y_i) \rightarrow 0$ as $y_i \uparrow \infty$. The boundary conditions have simplified the statement of results, but are not critical. If we assume only that $U_i(\cdot)$ is increasing, strictly concave, and continuously differentiable with $U_i'(0) < U_i'(\infty)$ then the Lyapunov function (5.1) remains strictly convex on the domain $\{b : U_i'(\infty) < b_i < U_i'(0), i \in \mathcal{I}\}$, has an interior minimum, and starting from any point $b(0)$ in this domain the trajectory $(b(t) : t \geq 0)$ converges to the point b achieving this minimum, which is the unique Nash equilibrium.*

6 General Assignments

In this section we consider how the assignment problem (2.1) can be generalized within our framework. Some extensions are immediate and straightforward. For example, we could allow the number of slots $L = L(\tau)$ to depend on the search type τ ; the pricing implementations of Section 2.2 do not require $L(\tau)$ to be constant over τ . In this Section we consider two further generalizations of practical importance.

6.1 More complex page layouts

Suppose the platform allows adverts of different sizes: for example, an advertiser may wish to offer an advert that occupies two adjacent slots. More generally adverts may vary in size, position, and include images and other media. So the platform may have a more complex set of possible page layouts than simply an ordered list of slots $1, 2, \dots, L$. Let $l \in \mathcal{L}$ describe a possible layout of the adverts for advertisers $i \in \mathcal{I}$. Let p_{il}^τ be the probability of a click-through to advertiser i under layout l . Then the generalization of the assignment problem (2.1) becomes

$$\begin{array}{ll} \text{Maximize} & \sum_{i \in \mathcal{I}} b_i p_{il}^\tau \\ \text{over} & l \in \mathcal{L}. \end{array}$$

Indeed, this formulation allows the click-through probabilities for an advert to depend not just on the advertiser and the position within the page, but also on which other adverts are shown on the page, provided only the probabilities p_{il}^τ can be estimated.

The complexity of this optimization problem depends on the design of the page layout through the structure of the set \mathcal{L} and may depend on any structural information on the probabilities p_{il}^τ , but for a variety of cases it will remain an assignment problem with an efficient solution. If $y_i(b)$ is again defined as the expected click-through rate for advertiser i from a bid vector b , and if it satisfies the monotonicity property, then Theorems 1 and 2 hold with identical proofs.

Example 3. *In an image-text auction, the platform may place on a page either an ordered set of text adverts (as described in Section 2) or a single image advert. As before advertiser i bids b_i , the marginal utility to advertiser i of an additional click-through; and now we suppose advertisers $i \in \mathcal{I}_{\text{text}}$ make available text adverts and advertisers $i \in \mathcal{I}_{\text{image}}$ make available image adverts, where $\mathcal{I} = \mathcal{I}_{\text{text}} \cup \mathcal{I}_{\text{image}}$ and an advertiser $i \in \mathcal{I}_{\text{text}} \cap \mathcal{I}_{\text{image}}$ makes available both a text and an image advert. Let the click-through probability on image advert i be p_i^τ for $i \in \mathcal{I}_{\text{image}}$, with click-through probabilities on text adverts as in Section 2.*

For this example the assignment problem is straightforward: the platform solves the earlier assignment problem (2.1) over advertisers $i \in \mathcal{I}_{\text{text}}$, and shows text adverts if the optimum achieved exceeds $\max_{i \in \mathcal{I}_{\text{image}}} b_i p_i^\tau$

and otherwise shows an image achieving this latter maximum. Similarly the calculation of the rebate is straightforward, with one further assignment problem to be solved for each click-through.

It is of course possible to construct assignment problems that are not as straightforward. For example, suppose that adverts are of different sizes, and the platform has a bound on the sum of the advert sizes shown. The assignment problem then includes as a special case the knapsack problem. In general the problem is NP-hard but it becomes computationally feasible if, for example, there are a limited number of possible advert sizes, as in the image-text auction above.

6.2 Controlling the number of slots

The platform may wish to limit the number of slots filled, if it judges the available adverts as not sufficiently interesting to searchers. Ultimately showing the wrong or poor quality adverts can cause searchers to move platform and so hurt long-term platform revenue.

Suppose the platform judges there is a benefit (positive or negative) q_{il}^τ to a searcher for an impression of the advert from advertiser i in slot l for a search of type τ , regardless of whether or not the searcher clicks on the advert. The system objective function (3.1a) then becomes

$$\sum_{i \in \mathcal{I}} U_i(y_i) + \mathbb{E}_\tau \left[\sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{L}} q_{il}^\tau x_{il}^\tau \right],$$

the assignment objective function (2.1a) becomes

$$\sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{L}} (b_i p_{il}^\tau + q_{il}^\tau) x_{il}^\tau,$$

and our results hold with minor amendments. In particular, equation (2.3) for the price function and equation (5.1) for the Lyapunov function are unaltered, although of course the functions $y_i(b)$ will now be defined in terms of solutions to the new assignment problem.

An important special case occurs when $q_{il}^\tau \equiv -R$, where R is a *reserve price*, but in this case we need to slightly perturb the set-up to ensure that Proposition 1 remains sufficient for the monotonicity property. Suppose that $q_{il}^\tau = q^\tau$ for all $i \in \mathcal{I}, l \in \mathcal{L}$ where $q^\tau = 0$ or $-R$ with probabilities ϵ and $1 - \epsilon$ respectively. (Formally, augment the space \mathcal{T} to carry a random variable q^τ that is independent of the click-through probabilities p_{il}^τ .) Then with probability $1 - \epsilon$ an advert will be shown in a slot only if its contribution to the objective function of the assignment problem, $b_i p_{il}^\tau$, is at least R . With probability ϵ a reserve is not applied: we add the possibility to ensure $y_i(b)$ is increasing even for small b_i .

Of course a reserve R may also have a favourable effect on the revenue received by the platform, Ostrovsky and Schwarz (2011); Bachrach et al. (2014). As an illustration, consider the generalized second price auction of Example 1. A reserve of R will reduce the number of slots filled if $R > b_L p_L^\tau$ and may increase the revenue, which can be calculated from expression (2.7). Nevertheless our framework is one of utility maximization: we assume the platform is trying to assure its long-term revenue by producing as much benefit as possible for its users, its advertisers and itself. There are, of course, several ways in which the platform could increase its own revenue within the utility maximization framework: in the absence of competition from other platforms it could for example charge an advertiser a fixed fee, less than the advertiser's consumer surplus, to participate.

As yet a further example of the flexibility of the framework, instead of a fixed reserve price we could allow an organic search result k to compete for a slot, with a positive benefit q_{kl}^τ , but with $b_k = 0$. Recent work has analyzed the trade-off in objectives between the platform and advertiser in sponsored search: Roberts et al. (2013) focus on ranking algorithms, trading off revenue against welfare, while Bachrach et al. (2014) also include the user as an additional stakeholder. Our framework aims to maximize the aggregate social welfare of the auction system, but it is noteworthy that this simple model of the benefit to a user of organic search results can be subsumed within our framework.

7 Platform-wide optimization

An advertiser may judge some types of click-through as more valuable than others. In this section we suppose that the platform allows an advertiser to express such preferences, by making distinct bids on different categories of search query. The challenge for the advertiser is to balance her bids across the range of categories offered to her by the platform.

Suppose the platform allows advertiser i to partition the type space \mathcal{T} into categories ($\mathcal{T}_{ik} : k \in \mathcal{K}_i$). The categories may be defined in terms of the keywords used in a search or any other feature of the search type, such as geographical area or broad classification of the user, that the platform is prepared to share with advertiser i . We assume the platform allows advertiser i to know the category of the search type τ , namely that $\tau \in \mathcal{T}_{ik}$,

but the platform knows more, namely τ . We suppose the platform may vary aspects of the auction, such as the number of advertising slots on the page or more generally the layout of the page, depending on the search type τ . For example, the platform may use the current screen size of the user to determine the page layout.

Let b_{ik} be the bid of advertiser i for click-throughs from category k , and let $b_i = (b_{ik} : k \in \mathcal{K}_i)$ and $b = (b_{ik} : i \in \mathcal{I}, k \in \mathcal{K}_i)$. Let y_{ik} be the click-through rate to advertiser i from searches in category $k \in \mathcal{K}_i$, and assume that the expected rate of payment by advertiser i for click-throughs from category $k \in \mathcal{K}_i$ is

$$\pi_{ik}(b)y_{ik}(b) = \int_0^{b_{ik}} \left(y_{ik}(b) - y_{ik}(b'_{ik}, b) \right) db'_{ik}, \quad i \in \mathcal{I}, k \in \mathcal{K}_i,$$

where (b'_{ik}, b) is the vector obtained from the vector b by replacing the component b_{ik} by b'_{ik} . This rate of payment can be achieved by either of the first two pricing implementations of Section 2.2: these implementations use the function $y_i^r(b)$ to determine the charge for a click-through, and so no difficulty is caused by the form of the auction depending upon the search type τ .

Let $y_i = (y_{ik} : k \in \mathcal{K}_i)$. If advertiser i 's utility $U_i(y_i)$ is simply a sum of utilities $U_{ik}(y_{ik})$ over the categories $k \in \mathcal{K}_i$ then this model is subsumed in the model treated in earlier sections: advertiser i can be represented by a collection of sub-advertisers, one for each category $k \in \mathcal{K}_i$, and the platform can set click-through probabilities to zero for sub-advertiser $k \in \mathcal{K}_i$ if $\tau \notin \mathcal{T}_{ik}$. But for more general utility functions we would expect that the bids $b_{ik}, k \in \mathcal{K}_i$, cannot be determined independently.

Suppose, then, that advertiser i 's utility $U_i(\cdot)$ is an increasing, strictly concave, continuously differentiable function of the vector $y_i = (y_{ik} : k \in \mathcal{K}_i)$. Assume that the partial derivative $\partial U_i / \partial y_{ik}$ decreases from ∞ to 0 as y_{ik} increases from 0 to ∞ , and that $b_{ik} \mapsto y_{ik}(b_{ik}, b)$ satisfies the monotonicity property.

Let

$$U_i^*(b_i) = \max_{y_i \geq 0} \left(U_i(y_i) - \sum_{k \in \mathcal{K}_i} b_{ik} y_{ik} \right),$$

the Legendre-Fenchel transform of $U_i(y_i)$, interpretable as the consumer surplus of advertiser i at prices b_i . Our conditions on U_i and its partial derivatives ensure there is a unique maximum, interior to the positive orthant, for any price vector b_i in the positive orthant. Let $(D_{ik}(b_i) : k \in \mathcal{K}_i)$ be the argument y_i that attains this maximum: it is the demand vector of advertiser i at prices b_i , and

$$\frac{\partial}{\partial b_{ik}} U_i^*(b_i) = -D_{ik}(b_i). \quad (7.1)$$

Then the question for advertiser i is how to balance her bids $(b_{ik} : k \in \mathcal{K}_i)$ over the categories \mathcal{K}_i that are of interest to her. The payoff to advertiser i arising from a vector of bids $b = (b_i : i \in \mathcal{I}) = (b_{ik} : i \in \mathcal{I}, k \in \mathcal{K}_i)$ is then

$$u_i(b) = U_i(y_i(b)) - \sum_{k \in \mathcal{K}_i} \pi_{ik}(b)y_{ik}(b),$$

and the condition for a Nash equilibrium is again (4.2) where now b_i is a vector. Paralleling the development of Section 4, the maximum of the payoff function $b_i \mapsto u_i(b_i, b_{-i})$ is attained when

$$\frac{\partial}{\partial b_{ik}} U_i^*(b_i) + y_{ik}(b) = 0, \quad k \in \mathcal{K}_i,$$

or equivalently $D_{ik}(b_i) = y_{ik}(b)$ for $k \in \mathcal{K}_i$, there is a unique Nash equilibrium, and these conditions also identify the unique system optimum.

Next suppose that for each $k \in \mathcal{K}_i$ advertiser i changes her bid $b_{ik}(t)$ smoothly (i.e., continuously and differentially) as a consequence of her observation of her current click-through rate $y_{ik}(b(t))$ so that

$$\frac{d}{dt} b_{ik}(t) \geq 0 \text{ according as } y_{ik}(b(t)) \leq D_{ik}(b_i(t)). \quad (7.2)$$

This is a dynamical system representation of advertiser i varying b_{ik} smoothly in order to increase or decrease her bid for keyword k according to whether the currently observed click-through rate $y_{ik}(t)$ is lower or higher than her demand at her current bid prices. Then trajectories converge to the solution of the system problem, by essentially the same Lyapunov argument as used to prove Theorem 2, as we now sketch.

Let

$$\mathcal{V}(b) = \sum_{i \in \mathcal{I}} U_i^*(b_i) + \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_i} b_{ik} y_{ik}(b).$$

Differentiating $\mathcal{V}(b(t))$ yields, from (7.1), Lemma 2 and (7.2),

$$\frac{d}{dt} \mathcal{V}(b(t)) = \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_i} \frac{\partial \mathcal{V}}{\partial b_{ik}} \frac{d}{dt} b_{ik}(t) = - \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_i} (D_{ik}(b_i(t)) - y_{ik}(b(t))) \frac{d}{dt} b_{ik}(t) \leq 0$$

where the inequality is strict unless $D_{ik}(b_i(t)) = y_{ik}(b(t))$ for $i \in \mathcal{I}, k \in \mathcal{K}_i$. But this holds if and only if y solves the system problem.

Remark 3. *Our approach to auction design separates the computational burden into a task that can be completed quickly for each page impression by the platform, and tasks that can be performed more slowly by individual advertisers or perhaps agents working on their behalf. The task for an advertiser is to assess the value to her of different forms of click-through. This may not be an easy task, but it is a task naturally assigned to the advertiser and is made simpler by requiring only local information in the region of the currently achieved click-through rates.*

Remark 4. *How finely should a search platform allow categories to be defined, and how finely should it divide its stream of queries across distinct auctions? Finer classifications will allow advertisers to communicate more precisely their valuations but excessive targeting may lead to thinner markets and to various forms of adverse selection. These trade-offs are discussed by Levin and Milgrom (2010), who argue that the degree of differentiation allowed, or conflation imposed, is an important aspect of the organization of well-functioning markets.*

In the current context, observe that advertiser i is forced to conflate her bid b_{ik} across multiple auctions, and in each of these the set of competing advertisers is likely to be different. Thus the design of the categories $(\mathcal{T}_{ik} : k \in \mathcal{K}_i, i \in \mathcal{I})$ provides ample opportunity to balance the degree of differentiation allowed to, or conflation imposed upon, advertisers by the platform.

We end this section with two examples which indicate the connections between our work and earlier important approaches in the traffic engineering and resource allocation literature.

Example 4. *Consider a platform with advertisers who prefer click-throughs that come from one geographical area rather than another, or from one set of keywords rather than another, simply because such click-throughs are more likely to convert into sales. Then advertiser i 's utility will be a univariate function*

$$U_i \left(\sum_{k \in \mathcal{K}_i} w_{ik} y_{ik} \right) \quad (7.3)$$

where we assume $U_i(\cdot)$ satisfies our earlier assumptions from Section 3 and where $w_i = (w_{ik} : k \in \mathcal{K}_i)$ account for the weight applied to each category by advertiser i .

Given that advertiser i declares a bid \tilde{b}_i and weights $\tilde{w}_i = (\tilde{w}_{ik} : k \in \mathcal{K}_i)$ (strategically and not necessarily equal to w_i), the platform may use the information contained in \tilde{w} as well as \tilde{b} and τ to solve the revised assignment problem, $\text{ASSIGNMENT}(\tau, \tilde{b}, \tilde{w})$, defined as problem (2.1) with the revised objective:

$$\text{Maximize} \quad \sum_{i \in \mathcal{I}} \tilde{b}_i \sum_{l \in \mathcal{L}} \tilde{w}_i^l p_{il}^T x_{il}^T,$$

where $\tilde{w}_i^l = \tilde{w}_{ik}$ for $\tau \in \mathcal{T}_{ik}$ and $k \in \mathcal{K}_i$. Write $b_{ik} = \tilde{b}_i \tilde{w}_{ik}$ and $b = (b_{ik} : k \in \mathcal{K}_i, i \in \mathcal{I})$. then the model formally reduces to that of this section. Essentially there are number of parallel auctions taking place with the search type τ determining the bids of advertisers: if $\tau \in \mathcal{T}_{ik}$ then the bid of advertiser i is b_{ik} .

The reduction replaces the vector function $U_i(y_{ik} : k \in \mathcal{K}_i)$ by the special case (7.3), a univariate function of a weighted sum. The utility (7.3) is concave but not strictly concave, and this introduces some minor technicalities and some simplifications. When strategic advertisers maximize over $(b_{ik} : k \in \mathcal{K}_i)$ their respective payoff functions

$$u_i(b) := U_i \left(\sum_{k \in \mathcal{K}_i} w_{ik} y_{ik}(b) \right) - \sum_{k \in \mathcal{K}_i} \pi_{ik}(b) y_{ik}(b)$$

the resulting Nash equilibrium is achieved under the necessary conditions

$$w_{ik} U_i' \left(\sum_{k \in \mathcal{K}_i} w_{ik} y_{ik}(b) \right) = b_{ik}, \quad \text{if } y_{ik}(b) > 0, \quad (7.4a)$$

$$w_{ik} U_i' \left(\sum_{k \in \mathcal{K}_i} w_{ik} y_{ik}(b) \right) \leq b_{ik}, \quad \text{if } y_{ik}(b) = 0, \quad (7.4b)$$

for $k \in \mathcal{K}_i$ and $i \in \mathcal{I}$. In particular, (7.4a) implies that if $y_{ik}(b) > 0$ then w_{ik}/b_{ik} does not depend on $k \in \mathcal{K}_i$: hence the weights intrinsic to advertiser i , $(w_{ik} : k \in \mathcal{K}_i)$, and the weights declared strategically by advertiser i to the platform, $(\tilde{w}_{ik} = b_{ik}/\tilde{b}_i : k \in \mathcal{K}_i)$, are in proportion wherever a positive click-through rate is received (this is the simplification achieved by the form (7.3)). So the auction mechanism achieves incentive compatibility: advertisers are encouraged to truthfully declare their intrinsic weights for categories for which they are competing.

We can interpret the system optimization as an infinitely large bipartite congestion game, an interpretation that parallels one of Vickrey's early motivations in transport pricing; in particular the conditions (7.4) parallel the conditions for a traffic equilibrium in a network, Wardrop (1952), Beckmann et al. (1956).

Example 5. An advertiser may have a budget constraint on what she can spend across different types of search, for example, in an advertising campaign. In this example we note a simple approach which captures a budget constraint within the framework of this section.

Suppose

$$U_i(y_i) = \frac{B_i}{q} \log \sum_{k \in \mathcal{K}_i} (w_{ik} y_{ik})^q$$

for $0 < q < 1$. Then

$$\frac{\partial U_i}{\partial y_{ik}} = \frac{B_i w_{ik}^q y_{ik}^{q-1}}{\sum_{j \in \mathcal{K}_i} (w_{ij} y_{ij})^q}$$

and so at the unique Nash equilibrium described earlier in this section, where $\partial U_i / \partial y_{ik} = b_{ik}$, the budget constraint

$$\sum_{k \in \mathcal{K}_i} b_{ik} y_{ik} = B_i$$

is automatically satisfied; note that the constraint is on the rate of bidding rather than expenditure, i.e., not taking into account rebates.

We require $q < 1$ to ensure the strict concavity of $U_i(\cdot)$. As $q \rightarrow 1$, maximizing $U_i(y_i)$ subject to the budget constraint approaches the problem of maximizing $\sum_{k \in \mathcal{K}_i} w_{ik} y_{ik}$ subject to the same budget constraint. In the special case where all advertisers are budget constrained we recover an important early model for the equilibrium price of goods for buyers with linear utilities, Fisher (1892); Eisenberg and Gale (1959).

Advertiser i will receive a stream of rebates, which may be delayed and will be noisy. Rebates will cause the total spent in a period to be less than the budget B_i , and a natural control response would be to spread the total rebate received in one time period over the budgets available for later time periods. Borgs et al. (2007) explore a natural bidding heuristic for a budget-constrained advertiser which readily extends to include delayed rebates.

8 Related Work

In this paper we have considered a problem where the social welfare of an auction system is optimized subject to the capacity constraints of that system. Social welfare optimization has long been an objective in the design of effective market mechanisms, Vickrey (1961). However, only in the recent literature have computationally efficient methods been considered for market and auction design, see Birnbaum et al. (2010), Jain and Vazirani (2007), Vazirani (2010). In the context of electronic commerce and specifically sponsored search auctions, these computational considerations are of critical importance given the increased diversity and competition associated with online advertising.

We have applied a decomposition approach to the task of optimizing advert allocation over the vast range of searches that can be conducted, and separated the task into sub-problems which can be implemented by each advertiser and on each search. The decompositions of interest are familiar and have been important in the context of communication network design, Srikant (2004); Kelly and Yudovina (2014). Tan and Srikant (2012) is a distinct approach using optimization decomposition ideas, but based instead on a queueing model of an on-line advert campaign and using connections to scheduling in wireless networks.

Strategic formulations of these optimization decompositions have been developed: in a simple model Johari and Tsitsiklis (2004) show a price of anarchy of 75% at a Nash equilibrium. Notably, a single parameter VCG mechanism to yield efficient allocation was considered by Maheswaran and Basar (2004) and subsequently generalized by Yang and Hajek (2007) and Johari and Tsitsiklis (2009). Here a parametrized surrogate utility is employed in the VCG mechanism, where the parameter is selected strategically. The message passed from the player to the mechanism is thus chosen from a reduced space. One part of our decomposition can be viewed as deriving a single parameter VCG mechanism with linear utilities; that is, despite allowing more general utility functions, the derived mechanism is *as if* each player had a linear utility function. Prior works have used strictly concave surrogate functions while in our approach linearisation is possible owing to the large search/constraint space employed. A linear VCG allocation is computationally straightforward (Leonard (1983), Bikhchandani et al. (2002)), but the crucial advantage of our linear framework is that – in addition to decomposing the objectives of advertisers and the platform – further decomposition over the search space is possible, leading to a practical mechanism. In particular, the mechanism can be implemented on each search instance: allocation and pricing both involve standard polynomial time algorithms per-search and per-click, respectively. In essence, we find a simple implementable auction mechanism that yields an efficient allocation of adverts across the entire search space.

Parametrized VCG mechanisms are examples of simplified mechanisms, where the set of messages available to report preferences is restricted. Milgrom (2010) has shown that the equilibria of a simplified mechanism are also equilibria of the unrestricted mechanism when a certain outcome closure property is satisfied. The closure property states that a bidder can make an optimal best response within the set of restricted bids whenever other

bidders' choices lie within the restricted set. As an example, the closure property can be applied to concave utility functions under the restriction of linearity, with the restricted bid communicating a tangent plane rather than the entire utility function.

A very common simplification applied in sponsored search auctions is conflation. For example advertisers may be required to make the same bid per click whatever the position of an advert on the page. If advertisers differentiate between positions beyond each position's observed click-through rate (for example, if click-throughs from lower positions are less or more valuable to the advertiser), then there may be a loss of social welfare from the restriction that a bidder must communicate a single parameter to a mechanism which is unaware of these positional effects. The question of whether VCG or GSP mechanisms with this restriction are sufficiently expressive to communicate the bidders' true values for positions is discussed in detail in Milgrom (2010) and Dütting et al. (2011). Aggarwal et al. (2007) discuss mechanisms which maintain an efficient equilibrium by allowing bidders to specify a minimum slot, in addition to their bid. Further recent discussions of simplified mechanisms for sponsored search auctions are Chawla and Hartline (2013); Hoy et al. (2013); Dütting et al. (2014).

The context in these papers is, in our terms, an auction for a single type of search. This context serves to illustrate an important theoretical question concerning simplified mechanisms. The diverse stochastic variability found in the sponsored search market (see Pin and Key (2011)) makes the assumption of a single type of search unrealistic. The framework we adopt is rather different: we presume, as described in Section 2, that the search platform knows more than the advertiser about the type of search being conducted, for example about the searcher, and that this information affects click-through probabilities. For the advertiser there is therefore a considerable further conflation: the same bid for a click-through is used over an entire category of search query as well as over different positions. The information asymmetry between the platform and the advertisers allows the platform to assign and price adverts using a per-search level of granularity on the search type, while the advertisers experience only average click-through rates over a diverse set of search types. Our framework is designed to model advertisers who differentiate the value of a click-through according to search categories, defined in terms of keywords and user characteristics, rather than advert position: see Section 7.

Borgs et al. (2007) gave an important early treatment of the dynamics of bid optimization, and emphasised the importance of equalizing the "return-on-investment" across keywords for budget-limited advertisers. These authors also used a continuous time formulation, noted the importance of random perturbations, proved convergence to a market equilibrium in the case of first price auctions and observed experimentally convergence in the case of second price auctions. Auctions were all single slot, and bids were assumed truthful. The additional contribution of Section 5 to this early work on dynamics is that we have established convergence for a wide class of continuous time dynamics representing advertisers' best response under our pricing mechanism.

More recent work on learning and bid optimization is reviewed by Tran-Thanh et al. (2014), who use the framework of a multi-armed bandit to devise policies that maximize the expected number of click-throughs in a given number of searches within a given budget. This stream of research typically uses no-regret learning, expressing convergence in terms of sub-linear temporal convergence. Iyer et al. (2011) use a mean field approach to treat agents who need to learn the value to them of a click-through. These are challenging problems, even for the sequence of single slot second price auctions treated in these papers. By comparison our approach uses a dynamical systems framework where fluid averages are controlled. We are also able to show convergence to a Nash equilibrium, rather than a correlated equilibrium. Our approach deliberately simplifies the modelling of the stochastic streams of click-throughs, which it represents with just their means, but is able to deal with multi-slot auctions and with streams of click-throughs arising from different keywords and categories of searcher. Nekipelov et al. (2015) review work on learning, focusing on a model of sponsored search auctions: their analysis of data from BingAds indicates that typical advertisers bid a significantly shaded version of their value, as would be expected in a generalized second price auction rather than a VCG auction.

As noted by Milgrom (2010), the most devastating objections to Vickrey pricing (Ausubel and Milgrom (2006)) apply only when bidders can buy multiple items, and have no force in sponsored search auctions where each bidder can acquire at most one position. Varian and Harris (2014) have recently argued that VCG mechanisms are of practical interest because they are flexible and extensible. For this reason, Facebook implements VCG² rather than the generalized second price auction currently used by Bing and Google. These considerations are particularly relevant for contextual advertisement, unordered page layouts, image-text adverts and image-video adverts. Such extensions are important and are of growing interest to online advertisement platforms; see Goel and Khani (2014) for a recent discussion.

9 Concluding Remarks

We describe a framework to capture the system architecture of Ad-auctions. The assignment problem must be solved rapidly, for each search, while an advertiser is primarily interested in aggregates over longer periods of

²<https://developers.facebook.com/docs/marketing-api/pacing> - downloaded on 20 August 2015.

time. The platform knows more about the type of a search query and thus more about click-through probabilities, while an advertiser knows more about the value to her of additional click-throughs and is incentivized to communicate this information via her bids. Thus we model in detail each random instance of the assignment problem, while we describe an advertiser’s strategic behaviour in terms of averages evolving in time. On a slow time-scale the platform may decide which search types to pool in distinct auctions, across which the advertisers will have different preferences they are able to communicate.

We have used sponsored search auctions as the motivation, and our model reflects current practice in sponsored search, where platforms such as BingAds or Google Adwords use a variant of the second price auction to solve the assignment and pricing problem for every search query, while advertisers alter bids on timescales measured in hours or days. The setting, allowing for a large, continuous range of search types and varying competition, greatly extends the scope of prior models which are typically limited to the auction of a single keyword amongst a static pool of advertisers.

We address the task of achieving efficiency over all a platform’s searches under a pay-per-click pricing model. Under the assumption of strategic advertisers, we showed that, with appropriate pricing, a Nash equilibrium exists for the advertisers which achieves welfare maximizing assignments. We gave an appealingly simple way to implement these prices: namely, by giving advertisers a rebate, constructed by solving a second assignment problem. The first assignment is implemented with low computation cost and the solution to the second assignment problem is not used for the allocation but only for pricing. Further this mechanism is found to be flexible and extends in a straightforward manner to various different page-layouts. Hence under the pay-per-click model, this mechanism shows potential to be adapted for use in current Ad-auctions.

References

- Gagan Aggarwal, Jon Feldman, and S Muthukrishnan. Bidding to the top: VCG and equilibria of position-based auctions. In *Approximation and Online Algorithms*, pages 15–28. Springer, 2007.
- Lawrence M Ausubel and Paul Milgrom. The lovely but lonely Vickrey auction. In Cramton P, Shoham Y, and Steinberg, editors, *Combinatorial Auctions*, pages 17–40. MIT Press, Cambridge, 2006.
- Yoram Bachrach, Sofia Ceppi, Ian A. Kash, Peter Key, and David Kurokawa. Optimising trade-offs among stakeholders in ad auctions. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC ’14. ACM, 2014.
- Martin Beckmann, CB McGuire, and Christopher B Winsten. *Studies in the Economics of Transportation*. Yale Univ. Press, New Haven, 1956.
- Dimitri P Bertsekas. The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of Operations Research*, 14(1):105–123, 1988.
- Sushil Bikhchandani, Sven de Vries, James Schummer, and Rakesh V Vohra. Linear programming and Vickrey auctions. In B Dietrich and RV Vohra, editors, *Mathematics of the Internet: E-auction and Markets. The IMA Volumes in Mathematics and Its Applications*, pages 75–116. Springer-Verlag, New York, 2002.
- Benjamin Birnbaum, Nikhil R. Devanur, and Lin Xiao. New convex programs and distributed algorithms for Fisher markets with linear and spending constraint utilities. Technical Report MSR-TR-2010-112, August 2010. URL <http://research.microsoft.com/apps/pubs/default.aspx?id=136327>.
- Christian Borgs, Jennifer Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th International Conference on World Wide Web*, WWW ’07, pages 531–540, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-654-7. doi: 10.1145/1242572.1242644. URL <http://doi.acm.org/10.1145/1242572.1242644>.
- J.M. Borwein and A.S. Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. CMS books in mathematics. Springer-Verlag New York, 2006.
- Shuchi Chawla and Jason D Hartline. Auctions with unique equilibria. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 181–196. ACM, 2013.
- Paul Dütting, Felix Fischer, and David C Parkes. Simplicity-expressiveness tradeoffs in mechanism design. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 341–350. ACM, 2011.
- Paul Dütting, Felix Fischer, and David C Parkes. Expressiveness and robustness of first-price position auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 57–74. ACM, 2014.

- 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.
- E. Eisenberg and D. Gale. Consensus of subjective probabilities: the pari-mutuel method. *The Annals of Mathematical Statistics*, 30:165–168, 1959.
- I. Fisher. Mathematical investigations in the theory of value and prices. *Transactions of the Connecticut Academy of Sciences and Arts*, 9:124, July 1892.
- Gagan Goel and Mohammad Reza Khani. Revenue monotone mechanisms for online advertising. In *Proceedings of the 23rd International Conference on World Wide Web, WWW '14*, pages 723–734, Republic and Canton of Geneva, Switzerland, 2014. International World Wide Web Conferences Steering Committee. ISBN 978-1-4503-2744-2. doi: 10.1145/2566486.2568000. URL <http://dx.doi.org/10.1145/2566486.2568000>.
- Darrell Hoy, Kamal Jain, and Christopher A Wilkens. A dynamic axiomatic approach to first-price auctions. *arXiv preprint arXiv:1304.7718*, 2013.
- Krishnamurthy Iyer, Ramesh Johari, and Mukund Sundararajan. Mean field equilibria of dynamic auctions with learning. *ACM SIGecom Exchanges*, 10(3):10–14, 2011.
- Kamal Jain and Vijay V. Vazirani. Eisenberg-Gale markets: algorithms and structural properties. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing, STOC '07*, pages 364–373, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-631-8. doi: <http://doi.acm.org/10.1145/1250790.1250845>. URL <http://doi.acm.org/10.1145/1250790.1250845>.
- Ramesh Johari and John N Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, 2004.
- Ramesh Johari and John N. Tsitsiklis. Efficiency of scalar-parameterized mechanisms. *Oper. Res.*, 57:823–839, July 2009.
- F Kelly. Fairness and stability of end-to-end congestion control. *European Journal of Control*, 9:159–176, 2003.
- F. Kelly and E. Yudovina. *Stochastic Networks*. Cambridge University Press, 2014.
- H.K. Khalil. *Nonlinear Systems*. Prentice Hall PTR, 2002.
- Harold W Kuhn. The Hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2): 83–97, 1955.
- Herman B Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *The Journal of Political Economy*, 91:461–479, 1983.
- Jonathan Levin and Paul Milgrom. Online advertising: heterogeneity and conflation in markt design. *American Economic Review*, 100(2):603–605, 2010.
- Rajiv T Maheswaran and Tamer Basar. Social welfare of selfish agents: motivating efficiency for divisible resources. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 2, pages 1550–1555. IEEE, 2004.
- Paul Milgrom. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior*, 70(1):62–70, 2010.
- P.R. Milgrom. *Putting Auction Theory to Work*. Churchill Lectures in Economics. Cambridge University Press, 2004. ISBN 9780521536721. URL <http://books.google.nl/books?id=AkeHTU7XW4kC>.
- Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- Denis Nekipelov, Vasilis Syrgkanis, and Eva Tardos. Econometrics for learning agents. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15*, pages 1–18, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3410-5. doi: 10.1145/2764468.2764522. URL <http://doi.acm.org/10.1145/2764468.2764522>.
- M. Ostrovsky and M. Schwarz. Reserve prices in Internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 59–60, 2011.
- F. Pin and P. Key. Stochastic variability in sponsored search auctions. In *Proceedings of the 12th ACM Conference on Electronic Commerce, EC'11*, pages 61–70, San Jose, California, 2011.

- Ben Roberts, Dinan Gunawardena, Ian A. Kash, and Peter Key. Ranking and tradeoffs in sponsored search auctions. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC '13*, pages 751–766. ACM, 2013. ISBN 978-1-4503-1962-1. doi: 10.1145/2482540.2482568. URL <http://doi.acm.org/10.1145/2482540.2482568>.
- R.T. Rockafellar. *Convex Analysis*. Convex Analysis. Princeton University Press, 1997. ISBN 9780691015866. URL <http://books.google.nl/books?id=1Ti0ka9bx3sC>.
- R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.
- Bo Tan and R. Srikant. Online advertisement, optimization and stochastic networks. *Automatic Control, IEEE Transactions on*, 57(11):2854–2868, Nov 2012. ISSN 0018-9286. doi: 10.1109/TAC.2012.2195810.
- Long Tran-Thanh, Lampros C. Stavrogiannis, Victor Naroditskiy, Valentin Robu, Nicholas R. Jennings, and Peter Key. Efficient regret bounds for online bid optimisation in budget-limited sponsored search auctions. In *30th Conference on Uncertainty in Artificial Intelligence*, pages 809–818, June 2014. URL <http://eprints.soton.ac.uk/365566/>.
- Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007. ISSN 0167-7187. doi: 10.1016/j.ijindorg.2006.10.002. URL <http://www.sciencedirect.com/science/article/pii/S0167718706001317>.
- Hal R Varian and Christopher Harris. The VCG auction in theory and practice. *The American Economic Review*, 104(5):442–445, 2014.
- Vijay V. Vazirani. Spending constraint utilities with applications to the adwords market. *Math. Oper. Res.*, 35:458–478, May 2010. doi: <http://dx.doi.org/10.1287/moor.1100.0450>. URL <http://dx.doi.org/10.1287/moor.1100.0450>.
- William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1): 8–37, 1961.
- John Glen Wardrop. Road paper. Some theoretical aspects of road traffic research. In *ICE Proceedings: Engineering Divisions*, volume 1, pages 325–362. Thomas Telford, 1952.
- S. Yang and B. Hajek. Revenue and stability of a mechanism for efficient allocation of a divisible good, 2006.
- Sichao Yang and B. Hajek. VCG-Kelly mechanisms for allocation of divisible goods: Adapting VCG mechanisms to one-dimensional signals. *Selected Areas in Communications, IEEE Journal on*, 25(6):1237–1243, August 2007. ISSN 0733-8716. doi: 10.1109/JSAC.2007.070817.

A Proof of Propositions 1 and Lemma 2

This section gives proofs of Proposition 1 and Lemma 2 concerning properties of the functions

$$y_i(b) = \mathbb{E}_\tau \sum_l p_{li}^\tau x_{li}^\tau(b), \quad \sum_i b_i y_i(b).$$

Proposition 1 requires the further technical lemma, Lemma 1, which give the Lipschitz continuity of a random point belonging to a polytope as we smoothly change the description of its facets. Lemma 2 employs the Envelope Theorem (Milgrom, 2004, Chap. 3), as is commonly applied in auction theory.

Lemma 1. 1) If U is a random vector uniformly distributed inside the unit sphere, $S_n = \{u \in \mathbb{R}^n : \|u\| \leq 1\}$, then there exist a constant K_1 such that for any two non-zero vectors $b, \tilde{b} \in \mathbb{R}^n \setminus \{0\}$

$$\mathbb{P}(b^T U \geq 0 > \tilde{b}^T U) \leq \frac{K_1}{\|b\| \wedge \|\tilde{b}\|} \|b - \tilde{b}\|.$$

2) If X is a random variable with density f_X continuous on its support \mathcal{P} , a polytope $\mathcal{P} \subset [-1, 1]^n$, then the function $\mathbb{P}(\mu_1^T X \geq 0, \dots, \mu_k^T X \geq 0)$ is Lipschitz continuous as a function of μ_1, \dots, μ_k provided $\|\mu_1\|, \dots, \|\mu_k\|$ are bounded away from zero.

Proof. 1) We give a geometric proof of the result. We assume, wlog, that $\|b\| \geq \|\tilde{b}\|$, and we let V_n be the volume of S . For every u satisfying $b^\top u \geq 0 > \tilde{b}^\top u$, there exists a $\theta \in [0, 1]$ such that $b^\top u + \theta(\tilde{b}^\top - b^\top)u = 0$. Let b_θ be the unit vector proportional to $b + \theta(\tilde{b} - b)$. So, $b_\theta^\top u = 0$. Or, in other words, if $b^\top u + \theta(\tilde{b}^\top - b^\top)u = 0$ then u belongs to a cross section of the sphere $\{u' \in S : b_\theta u' = 0\}$ for some b_θ proportional to $b + \theta(\tilde{b} - b)$. Thus the volume of $\{u : b^\top u \geq 0 > \tilde{b}^\top u\}$ is overestimated by the area of the sets $\{u' \in S : b_\theta u' = 0\}$ multiplied by the change in the normal vector b_θ .

With this in mind we note three facts: 1) Each cross section $\{u \in S : b_\theta^\top u = 0\}$ has the same volume V_{n-1} in its $n-1$ dimensional subspace; 2) The path $\mathcal{P} = \{b_\theta : \theta \in [0, 1]\}$ is a circular path starting at $b/\|b\|$ and ending at $\tilde{b}/\|\tilde{b}\|$, and thus has length bounded above by the terms

$$2\pi \left\| \frac{b}{\|b\|} - \frac{\tilde{b}}{\|\tilde{b}\|} \right\| \leq \frac{2\pi}{\|\tilde{b}\|} \|b - \tilde{b}\|;$$

and, 3) $\{u \in S : b^\top u \geq 0 > \tilde{b}^\top u\} = \{u \in S : b_\theta^\top u = 0, \theta \in [0, 1]\}$. Thus, we see we can bound the probability $\mathbb{P}(b^\top U \geq 0 > \tilde{b}^\top U)$ by the length of the path \mathcal{P} times the volume of cross sections $\{u \in S : b_\theta^\top u = 0\}$. In other words,

$$\mathbb{P}(b^\top U \geq 0 > \tilde{b}^\top U) \leq \frac{2\pi V_{n-1}}{\|\tilde{b}\|} \|b - \tilde{b}\|,$$

as required.

2) A function which is componentwise Lipschitz continuous is Lipschitz continuous. So, without loss of generality, we prove that the first component of our function is Lipschitz continuous. Observe

$$\begin{aligned} & \left| \mathbb{P}(\mu_1^\top X \geq 0, \mu_2^\top X \geq 0, \dots, \mu_k^\top X \geq 0) - \mathbb{P}(\tilde{\mu}_1^\top X \geq 0, \mu_2^\top X \geq 0, \dots, \mu_k^\top X \geq 0) \right| \\ &= \left| \mathbb{P}(\mu_1^\top X \geq 0 > \tilde{\mu}_1^\top X, \mu_2^\top X \geq 0, \dots, \mu_k^\top X \geq 0) - \mathbb{P}(\tilde{\mu}_1^\top X \geq 0 > \mu_1^\top X, \mu_2^\top X \geq 0, \dots, \mu_k^\top X \geq 0) \right| \\ &\leq \mathbb{P}(\mu_1^\top X \geq 0 > \tilde{\mu}_1^\top X) + \mathbb{P}(\tilde{\mu}_1^\top X \geq 0 > \mu_1^\top X). \end{aligned} \quad (\text{A.1})$$

Also since f is a continuous density on click-through probabilities $\tilde{\mathcal{P}}$, it is bounded by a constant. So, we can bound the above probabilities with uniform random variables:

$$\mathbb{P}(\mu^\top X \geq 0 > \tilde{\mu}^\top X) \leq K_2 \mathbb{P}(\mu^\top U \geq 0 > \tilde{\mu}^\top U)$$

for a constant K_2 and for U a uniform random variable on the unit sphere in \mathbb{R}^n . Now applying part 1) of this Lemma

$$\mathbb{P}(\mu^\top X \geq 0 > \tilde{\mu}^\top X) \leq \frac{K_1 K_2}{\|\mu\| \wedge \|\tilde{\mu}\|} \|\mu - \tilde{\mu}\| \leq \frac{K_1 K_2}{K_3} \|\mu - \tilde{\mu}\|.$$

where K_3 is the constant by which μ and $\tilde{\mu}$ are bounded away from zero. Thus, applying this inequality to (A.1), we have that $\mathbb{P}(\mu_1^\top X \geq 0, \mu_2^\top X \geq 0, \dots, \mu_k^\top X \geq 0)$ is Lipschitz continuous in its first component and thus is Lipschitz continuous. \square

The previous lemmas suggest that provided there is a certain amount of variability in p_{ij}^τ then we can expect the average performance of an advertiser to be a continuous function of the declared prices b .

Proof of Proposition 1. First we argue the continuity of $b_i \mapsto y_i(b_i, b_{-i})$ and then argue that it is strictly increasing and positive. We let \mathcal{S} index the assignments that can be scheduled from \mathcal{I} to \mathcal{L} . Notice, provide there is a unique maximal assignment,

$$\begin{aligned} x_{il}^\tau(b) &= \sum_{\pi \in \mathcal{S} : \pi(i)=l} \mathbf{I} \left[\sum_k b_k p_{k\pi}^\tau \geq \sum_k b_k p_{k\tilde{\pi}}^\tau, \forall \tilde{\pi} \neq \pi \right] \\ &= \sum_{\pi \in \mathcal{S} : \pi(i)=l} \prod_{\tilde{\pi} \neq \pi} \mathbf{I} \left[\sum_k b_k p_{k\pi}^\tau \geq \sum_k b_k p_{k\tilde{\pi}}^\tau \right]. \end{aligned} \quad (\text{A.2})$$

Here \mathbf{I} is the indicator function. Notice, since \mathbb{P}_τ admits a density, $f(p^\tau)$, then with probability one there is a unique maximizer to the problem ASSIGNMENT(τ, b). So the equality (A.2) holds almost surely for all $b > 0$.

For two assignments π and $\tilde{\pi}$, we define the vector

$$\mu_{\pi\tilde{\pi}} := (b_i \mathbf{I}[\pi(i) = l] - b_i \mathbf{I}[\tilde{\pi}(i) = l]) : i \in \mathcal{I}, l \in \mathcal{L}.$$

Notice for any two distinct permutations, the non-zero components of $\mathbf{I}[\tilde{\pi}(i) = l]$ are distinct. So the vectors $\mu_{\pi\tilde{\pi}}$ are distinct and non-zero over $\tilde{\pi} \neq \pi$. Since the maximal assignment is almost surely unique, we have

$$x_{il}(b) = \sum_{\pi \in \mathcal{S}: \pi(i)=l} \mathbb{E} \left[\prod_{\tilde{\pi} \neq \pi} \mathbf{I} \left[\sum_k b_k p_{k\pi(k)}^\tau \geq \sum_k b_k p_{k\tilde{\pi}(k)}^\tau \right] \right] = \sum_{\pi \in \mathcal{S}: \pi(i)=l} \mathbb{P}(\mu_{\pi\tilde{\pi}}^\top p \geq 0, \forall \tilde{\pi} \neq \pi). \quad (\text{A.3})$$

Thus if the function $\mathbb{P}(\mu_{\pi\tilde{\pi}}^\top p \geq 0, \forall \tilde{\pi} \neq \pi)$ is Lipschitz continuous then we have same properties for functions $x_{jl}(b)$. The Lipschitz continuity of $\mathbb{P}(\mu_{\pi\tilde{\pi}}^\top p \geq 0, \forall \tilde{\pi} \neq \pi)$ is proven in Lemma 1. This implies the Lipschitz property for $x_{il}(b)$ with $b > 0$ and since y_i is a finite sum of these terms the same continuity holds for $b_i \mapsto y_i(b_i, b_{-i})$, with $b = (b_i, b_{-i}) > 0$. Further, continuity at $b_i = 0$ is also ensured by bounded convergence: there are greater than $|\mathcal{L}|$ positive bids occur in b_{-i} the assignment of these must eventually outweighs the assignment of i as $b_i \searrow 0$. In other words, (A.2) goes to zero point-wise as $b_i \searrow 0$. Thus bounded convergence applies to (A.3) and, also, $y_i(b)$ which implies $y_i(b_i, b_{-i}) \rightarrow 0$ as $b_i \searrow 0$.

We now prove that the function $b_i \mapsto y_i(b_i, b_{-i})$ is strictly increasing for $b_{-i} \neq 0$. First we show that it is increasing. Since $y_i(b) = \mathbb{E}_\tau y_i^\tau(b)$ (2.2), if we can prove $y_i^\tau(b)$ is increasing then so is $y_i(b)$. Further note

$$\sum_{i \in \mathcal{I}} b_i y_i^\tau(b) = \sum_{i \in \mathcal{I}, l \in \mathcal{L}} b_i p_{il}^\tau x_{il}^\tau(b)$$

which is the optimal objective for the assignment problem (2.1).

Define b' with $b'_i < b_i$ and $b'_j = b_j$ for each $j \neq i$. We now proceed by contradiction. Suppose that $y_i(b') > y_i(b)$, then the following equalities and inequalities hold

$$\begin{aligned} \sum_{j \in \mathcal{I}} b_j y_j^\tau(b) &= (b_i - b'_i) y_i^\tau(b) + \sum_{j \in \mathcal{J}} b'_j y_j^\tau(b) \\ &\leq (b_i - b'_i) y_i^\tau(b) + \sum_{j \in \mathcal{J}} b'_j y_j^\tau(b') \\ &< (b_i - b'_i) y_i^\tau(b') + \sum_{j \in \mathcal{J}} b'_j y_j^\tau(b') = \sum_{j \in \mathcal{J}} b_j y_j^\tau(b'). \end{aligned}$$

Here the first equality holds by the optimality of $y^\tau(b')$ and the second holds by assumption. But notice the resulting equality above contradicts the optimality of $y^\tau(b)$. Thus by contradiction, $y_i^\tau(b)$ is increasing in b_i and, after taking expectations, so is $y_i(b)$.

We now prove that $b_i \mapsto y_i(b)$ is strictly increasing. Let b' be such that $b'_i > b_i$ and $b'_j = b_j$ for all $j \neq i$. The result proceeds by showing that

$$\mathbb{P}(y_i^\tau(b') > y_i^\tau(b) | E) > 0$$

where we condition on an event E with non-zero probability. Notice, after taking expectations, this implies that $y_i(b') > y_i(b)$.

Now since $f(p)$ is positive on a region containing the origin, $f(p)$ stochastically dominates a uniform random variable on the set of increasing click-through rates, $\tilde{\mathcal{P}} \cap [0, \epsilon]^{\mathcal{I} \times \mathcal{L}}$, for some ϵ . Thus it is sufficient to prove the result for $u = (u_{il} : i \in \mathcal{I}, l \in \mathcal{L})$ uniform on $\tilde{\mathcal{P}} \cap [0, \epsilon]^{\mathcal{I} \times \mathcal{L}}$. Now, for instance, there is positive probability that advertiser i and j , with $b_j > 0$, compete exclusively over the top two slots, $l = 1, 2$. This occurs, for instance, when i and j have click-through rate over $\epsilon/2$ and all other advertisers have expected revenue that is half of the lower bound revenue of i and j , namely, the event

$$E := \left\{ \min_{\substack{k=i,j \\ l \in \mathcal{L}}} u_{kl} \geq \frac{\epsilon}{2}, \quad 2 \max_{\substack{k \neq i,j \\ l \in \mathcal{L}}} \{b_k u_{kl}\} \leq \frac{\epsilon}{2} \min\{b_i, b_j\} \right\}.$$

This event has positive probability and then only i and j can appear on the top two slots on this event.

Given this event, advertiser i achieves the top position with bid b'_i and the second position with bid b_i on the condition

$$b'_i(u_{i1} - u_{i2}) > b_j(u_{j1} - u_{j2}) > b_i(u_{i1} - u_{i2}).$$

Since, after conditioning on E , $u_{i1}, u_{i2}, u_{j1}, u_{j2}$ remain independent and uniformly distributed (on the set $\tilde{\mathcal{P}} \cap \{u_{i1}, u_{i2}, u_{j1}, u_{j2} \geq \epsilon/2\}$), it is a straightforward calculation that

$$\mathbb{P}(b'_i(u_{i1} - u_{i2}) > b_j(u_{j1} - u_{j2}) > b_i(u_{i1} - u_{i2}) | E) > 0.$$

Since u_{i1} , the value of y_i^τ achieved by b'_i on E , is strictly bigger than u_{i2} , the value of y_i^τ achieved by b_i on E , the above inequality implies

$$\mathbb{P}(y_i^\tau(b') > y_i^\tau(b) | E) > 0,$$

and thus $y_i(b) < y_i(b')$, as required. Further, note that this argument implies the required property that $y_i(b_i, b_{-i}) > 0$ for $b_i > 0$. \square

Lemma 2. The function $b \mapsto \sum_{i \in \mathcal{I}} b_i y_i(b)$ is convex and continuously differentiable for $b \neq 0$; further,

$$\frac{d}{db_i} \left\{ \sum_{i' \in \mathcal{I}} b_{i'} y_{i'}(b) \right\} = y_i(b).$$

and

$$\lim_{\|b\| \rightarrow \infty} \sum_{i \in \mathcal{I}} b_i y_i(b) = \infty. \quad (\text{A.4})$$

Proof. The optimal value of the assignment problem (2.1) is convex as a function of b , since it is the supremum of a set of linear functions. Thus the function $b \mapsto \sum_{i \in \mathcal{I}} b_i y_i(b)$, a linear combination of convex functions, is also convex. Further $\sum_{i \in \mathcal{I}} b_i y_i(\tilde{b})$ is a supporting hyperplane at the point \tilde{b} . Differentiability can be shown to follow from the continuity of $y(\tilde{b})$ as a function of \tilde{b} , and we next give a detailed proof following the Envelope Theorem (Milgrom, 2004, Chap. 3).

Since by definition, $x^\tau(b)$ is optimal for the assignment problem, we have that

$$\sum_{i \in \mathcal{I}} b_i y_i(b) = \mathbb{E} \left[\sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{L}} b_i p_{il}^\tau x_{il}^\tau(b) \right] = \mathbb{E} \left[\max_{x^\tau \in \mathcal{S}} \sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{L}} b_i p_{il}^\tau x_{il}^\tau \right],$$

Letting $b^h = b + e_i h$, where e_i is the i th unit vector in $\mathbb{R}^{\mathcal{I}}$ and $h > 0$ (a symmetric argument holds for $h < 0$). We see that the partial derivative with respect to b_i is lower-bounded

$$\begin{aligned} \sum_{i' \in \mathcal{I}} \frac{b_{i'}^h y_{i'}(b^h) - b_{i'} y_{i'}(b)}{h} &\geq \frac{1}{h} \left\{ \mathbb{E} \left[\sum_{i' \in \mathcal{I}} \sum_{l \in \mathcal{L}} b_{i'}^h p_{i'l}^\tau x_{i'l}^\tau(b) \right] - \mathbb{E} \left[\sum_{i' \in \mathcal{I}} \sum_{l \in \mathcal{L}} b_{i'} p_{i'l}^\tau x_{i'l}^\tau(b) \right] \right\} \\ &= \mathbb{E} \left[\sum_{l \in \mathcal{L}} p_{il}^\tau x_{il}^\tau(b) \right] = y_i(b). \end{aligned} \quad (\text{A.5})$$

The inequality above holds because $x^\tau(b)$ is suboptimal for the parameter choice b^h . By the same argument, applied to $b_i y_i(b)$ instead of $b_{i'} y_{i'}(b)$ in (A.5), we also have that

$$\sum_{i' \in \mathcal{I}} \frac{b_{i'}^h y_{i'}(b^h) - b_{i'} y_{i'}(b)}{h} \leq y_i(b^h).$$

By the continuity of $y_i(b)$, letting $h \rightarrow 0$ gives that

$$\frac{d}{db_i} \sum_{i' \in \mathcal{I}} b_{i'} y_{i'}(b) = y_i(b),$$

as required.

Since $y_i(b_i, b_{-i})$ is non-zero for $b_i > 0$ it follows that

$$\min_{\|b\|=1} \sum_{i \in \mathcal{I}} b_i y_i(b) > 0,$$

and consequently, letting $\|b\| \rightarrow \infty$, we see that (A.4) holds. \square

B Proof of Propositions 2 and 3

Proof of Proposition 2. A Lagrangian of the system problem (3.1a), (3.1b) can be written as follows

$$L_{sys}(x, y; b) = \sum_{i \in \mathcal{I}} U_i(y_i) + \sum_{i \in \mathcal{I}} b_i \mathbb{E}_\tau \left[\sum_{l \in \mathcal{L}} p_{il}^\tau x_{il}^\tau - y_i \right]. \quad (\text{B.1})$$

Note that we intentionally omit the scheduling constraints from our Lagrangian, and therefore we must maximize subject to these constraints, (3.1c-3.1d), when optimizing our Lagrangian.

Let \mathcal{A} be the set of variables $x = (x^\tau \in \mathcal{S} : \tau \in \mathcal{T})$ satisfying the assignment constraints (3.1c-3.1e). We see that our Lagrangian problem is separable in the following sense

$$\max_{\substack{x \in \mathcal{A} \\ y, z \in \mathbb{R}_+^{\mathcal{I}}}} L_{sys}(x, y; b) = \sum_{i \in \mathcal{I}} \max_{y_i \geq 0} \{U_i(y_i) - b_i y_i\} \quad (\text{B.2a})$$

$$+ \mathbb{E}_\tau \left[\max_{x^\tau \in \mathcal{S}} \sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{L}} b_i p_{il}^\tau x_{il}^\tau \right] \quad (\text{B.2b})$$

Here \mathcal{S} denotes the set of $x' \in \mathbb{R}_+^{\mathcal{I} \times \mathcal{L}}$ such that for each $i \in \mathcal{I}$ and $l \in \mathcal{L}$

$$\sum_{l' \in \mathcal{L}} x'_{il'} \leq 1, \text{ and } \sum_{i' \in \mathcal{I}} x'_{i'l} \leq 1.$$

We now show that solutions \tilde{x} , \tilde{y} and \tilde{b} satisfying the Conditions A and B of our Proposition are optimal for the Lagrangian (B.2a) and (B.2b) when $b = \tilde{b}$.

Firstly, suppose Conditions A and B are satisfied. Assuming Condition A, the following is a straightforward application of Fenchel duality. If \tilde{b}_i a solution to the optimization

$$\text{minimize } U_i^*(b_i) + b_i \tilde{y}_i \text{ over } b_i \geq 0,$$

then, under our expression (3.4) for U^* , the solution is achieved when $D_i(\tilde{b}_i) = \tilde{y}_i$ or equivalently when $U_i'(\tilde{y}_i) = \tilde{b}_i$. Thus it is clear that \tilde{y}_i solves the optimization

$$\max_{y_i \geq 0} \{U_i(y_i) - \tilde{b}_i y_i\}.$$

Hence if Condition A is satisfied, then \tilde{y}_i optimizes (B.2a) when we choose $b_i = \tilde{b}_i$.

Secondly, if \tilde{x}^τ solves ASSIGNMENT(τ, \tilde{b}) for each τ , because each maximization inside the expectation (B.2b) is an assignment problem, then (B.2b) is maximized by \tilde{x} when we take $b = \tilde{b}$.

These two conditions, Condition A and B, show that the Lagrangian (B.1) is maximized by \tilde{x} and \tilde{y} with Lagrange multipliers \tilde{b} . In addition, \tilde{x} and \tilde{y} are feasible for the system optimization (3.1) and hence we have a feasible optimal solution for this Lagrangian problem. But as we demonstrate in Proposition 5 below, Lagrangian sufficiency still holds for the system problem (3.1) – despite the infinite number of constraints. Therefore we have shown a solution to Conditions A and B is optimal for the system problem.

Conversely, we know that strong duality holds for the system optimization (3.1) – even with the infinite number of constraints for this optimization (see Theorem 3 in the Appendix for a proof). Hence there exists a vector \tilde{b} such that an optimal solution to the system problem is also an optimal solution to the Lagrangian problem when we chose Lagrange multipliers \tilde{b} . Thus, an optimal solution to the SYSTEM($U, \mathcal{I}, \mathbb{P}_\tau$) must optimize (B.2a) and (B.2b), and as discussed these solutions correspond to Conditions A and B. In other words, an optimal solution to the system problem satisfies Conditions A and B with this choice of \tilde{b} . \square

Proof of Proposition 3. a) From Theorem 2, the Lagrangian of the system problem can be written as follows

$$L_{sys}(x, y; b) = \sum_{i \in \mathcal{I}} U_i(y_i) + \sum_{i \in \mathcal{I}} b_i \mathbb{E}_\tau \left[\sum_{l \in \mathcal{L}} p_{il}^\tau x_{il}^\tau - y_i \right]. \quad (\text{B.3})$$

Recall from (B.2), this Lagrangian is separable and is maximized as

$$\begin{aligned} \max_{\substack{\tilde{x} \in \mathcal{A} \\ y, z \in \mathbb{R}_+^{\mathcal{I}}}} L_{sys}(x, y; b) &= \sum_{i \in \mathcal{I}} \max_{y_i \geq 0} \{U_i(y_i) - b_i y_i\} + \mathbb{E}_\tau \left[\max_{x^\tau \in \mathcal{S}} \sum_{i \in \mathcal{I}} \sum_{l \in \mathcal{L}} b_i p_{il}^\tau x_{il}^\tau \right] \\ &= \sum_{i \in \mathcal{I}} U_i^*(b_i) + \mathbb{E}_\tau \left[\sum_{i \in \mathcal{I}} b_i \sum_{l \in \mathcal{L}} p_{il}^\tau x_{il}^\tau(b) \right] = \sum_{i \in \mathcal{I}} (U_i^*(b_i) + b_i y_i(b)). \end{aligned}$$

In the second equality above, we rearrange the assignment optimization in terms of the click-through rate of each advertiser, $y_i(b)$.

Thus the dual of this optimization problem is as required:

$$\text{Minimize } \sum_{i \in \mathcal{I}} [U_i^*(b_i) + b_i y_i(b)] \text{ over } b_i \geq 0, \quad i \in \mathcal{I}.$$

We analyze this dual problem. We first show that optimization (3.5) is minimized when $0 < b_i < \infty$ for each $i \in \mathcal{I}$. We consider the function

$$\sum_{i \in \mathcal{I}} b_i y_i(b).$$

With the technical lemma, Lemma 2, we see that this function is continuous and, for $b_{-i} \neq 0$, differentiable with i th partial derivative given by the continuous function $y_i(b)$. Further it satisfies (A.4). Thus since $U_i^*(b_i)$ is a positive function, we see that the dual minimization (3.5) must be achieved by a finite solution b^* . In addition, by definition $D_i(b) = -(U_i^*)'(b) = (U_i')^{-1}(b)$, and so the objective of the dual is continuously differentiable for $b > 0$. Since $D_i(0) = \infty$, the minimum of the dual problem (3.5) must be achieved by $b_i^* > 0$ for each $i \in \mathcal{I}$.

Now, as the objective of (3.5) is continuously differentiable for b strictly positive, it is minimized iff for each $i \in \mathcal{I}$

$$\frac{dU_i^*}{db_i}(b_i^*) + y_i(b^*) = 0.$$

Finally, since each function U_i^* is strictly convex, dual objective is strictly convex and so the above minimizer is unique.

b) For the Lagrangian for the system problem, (B.3), Strong Duality holds by Theorem 3. So, there exist Lagrange multipliers b^* , such that

$$\sum_{i \in \mathcal{I}} [U_i^*(b_i^*) + b_i^* y_i(b^*)] = \max_{\substack{x \in \mathcal{A} \\ y \in \mathbb{R}_+^{\mathcal{I}}}} L_{sys}(x, y; b^*) = \max_{\substack{x \in \mathcal{A} \\ y \in \mathbb{R}_+^{\mathcal{I}}}} \sum_{i \in \mathcal{I}} U_i(y_i)$$

where there are feasible vectors x^* , y^* achieving the optimum of both maximizations above. By weak duality it is clear that b^* must be optimal for the dual problem (3.5). Further, since x^* optimizes the Lagrangian L_{sys} with Lagrange multipliers b^* , it solves the assignment problem, $x^{*\tau} = x^\tau(b^*)$. \square

C Lagrangian Optimization

In this paper, we consider optimization problems that have a potentially infinite number of constraints, in particular, for the system-wide optimization (3.1). Thus it is not immediately clear that the Lagrangian approach – ordinarily applied with a finite number of constraints – immediately applies to our setting. We demonstrate that certain principle results, namely weak duality, the Lagrangian Sufficiency and strong duality, apply to our setting. These technical lemmas supplement proofs in Propositions 2 and 3.

We consider an optimization of the form

$$\text{Maximize} \quad g(y) \quad (\text{C.1a})$$

$$\text{subject to} \quad y_i \leq \mathbb{E}_\mu[x_i], \quad i = 1, \dots, n, \quad (\text{C.1b})$$

$$f_j(x(\tau)) \leq c_j, \quad \tau \in \mathcal{T}, \quad j = 1, \dots, m, \quad (\text{C.1c})$$

$$\text{over} \quad y \in \mathbb{R}^n, \quad x \in \mathcal{B}(\mathcal{T}, \mathbb{R}^n). \quad (\text{C.1d})$$

In the above optimization, we consider probability space $(\mathcal{T}, \mathbb{P}_\mu)$ and measurable random variable $x : \mathcal{T} \rightarrow \mathbb{R}^n$. We let $\mathcal{B}(\mathcal{T}, \mathbb{R}^n)$ index the set of Borel measurable functions from \mathcal{T} to \mathbb{R}^n . We assume that $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is a concave function and that $f_j : \mathbb{R}^n \rightarrow \mathbb{R}$ is a convex function, for each $j = 1, \dots, m$. We assume the solution to this optimization is bounded above.

Although there are an infinite number of constraints in this optimization, we can define a Lagrangian for this optimization as follows

$$L(x, y, z; b) = g(y) + \sum_{i=1}^n b_i \mathbb{E}_\mu[x_i - y_i - z_i].$$

Here the Lagrange multipliers b_i , $i = 1, \dots, n$, can be assumed to be positive, slack variables z_i are added for each constraint (C.1b) and the optimization of the Lagrangian is taken over y_i real, z_i positive and real, and x_i a Borel measurable random variable for $i \in \mathcal{I}$. We let \mathcal{F} be the set of (x, y) feasible for the optimization (C.1).

Weak duality and Lagrangian Sufficiency both hold for this Lagrangian problem.

Proposition 5 (Weak Duality).

a) [Weak Duality] For g^* the optimal value of the optimization (C.1),

$$\sup_{\substack{y \in \mathbb{R}^n, \\ x \in \mathcal{B}(\mathcal{T}, \mathbb{R}^n)}} L(x, y, z; b) \geq g^*.$$

b) [Lagrangian Sufficiency] If, given some b , there exists $x^* \in \mathcal{B}(\mathcal{T}, \mathbb{R}^n)$ and $y^*, z^* \in \mathbb{R}^n$ that are feasible for the optimization (C.1) and maximize the Lagrangian $L(x, y, z; b)$ with $z_i^* := y_i^* - \mathbb{E}_\mu x_i^*$ then x^*, y^*, z^* is optimal for (C.1).

Proof. a) Because \mathcal{F} is a subset of $\mathcal{B}(\mathcal{T}, \mathbb{R}^n) \times \mathbb{R}^n$, we have

$$\sup_{\substack{y \in \mathbb{R}^n, z \in \mathbb{R}_+^n, \\ x \in \mathcal{B}(\mathcal{T}, \mathbb{R}^n)}} L(x, y, z; b) \geq \sup_{\substack{(x, y) \in \mathcal{F} \\ z \in \mathbb{R}_+^n}} L(x, y, z; b) = g^*.$$

This proves weak duality.

b) Now applying this inequality, if a feasible solution optimizes the Lagrangian, then

$$g(y^*) = L(x^*, y^*, z^*; b) = \sup_{\substack{y \in \mathbb{R}^n, z \in \mathbb{R}_+^n, \\ x \in \mathcal{B}(\mathcal{T}, \mathbb{R}^n)}} L(x, y, z; b) \geq g^*.$$

Thus, (x^*, y^*) is optimal for (C.1). \square

For $z \in \mathbb{R}^{\mathcal{I}}$, we use $\mathcal{F}(z)$ to denote the set of (x, y) satisfying constraints (C.1c-C.1d) and satisfying constraints

$$z_i + y_i \leq \mathbb{E}_\mu[x_i], \quad i = 1, \dots, n.$$

Note, $\mathcal{F} = \mathcal{F}(0)$. We now show that there exists a Lagrange multiplier b^* where the optimized Lagrangian function also optimizes (C.1).

Theorem 3 (Strong Duality). *There exists a $b^* \in \mathbb{R}_+^n$ such that*

$$\max_{(x, y) \in \mathcal{F}} g(y) = \max_{\substack{y \in \mathbb{R}^n \\ x \in \mathcal{B}(\mathcal{T}, \mathbb{R}^n)}} g(y) + \sum_{i \in \mathcal{I}} b_i^* \mathbb{E}[x_i - y_i]. \quad (\text{C.2})$$

In particular, if there exist $(x^, y^*) \in \mathcal{F}$ maximizing (C.1) then it maximizes (C.2).*

Proof. Firstly, since $\mathcal{F} \subset \mathcal{B}(\mathcal{T}, \mathbb{R}^n) \times \mathbb{R}^n$, we proved the weak duality expression

$$\max_{(x, y) \in \mathcal{F}} g(y) = \max_{(x, y) \in \mathcal{F}} g(y) + \sum_{i \in \mathcal{I}} b_i^* \mathbb{E}[x_i - y_i] \leq \max_{\substack{y \in \mathbb{R}^{\mathcal{I}} \\ x \in \mathcal{B}(\mathcal{T}, \mathbb{R}^n)}} g(y) + \sum_{i \in \mathcal{I}} b_i^* \mathbb{E}[x_i - y_i]. \quad (\text{C.3})$$

So it remains to show the reverse inequality. We consider the following set

$$\mathcal{C} = \{(z, \gamma) \in \mathbb{R}^{\mathcal{I}} \times \mathbb{R} : \text{there exists } (x, y) \in \mathcal{F}(z) \text{ with } g(y) \geq \gamma\}.$$

We claim that \mathcal{C} is convex. Take $(z^0, \gamma^0), (z^1, \gamma^1) \in \mathcal{C}$ and take $(x^0, y^0) \in \mathcal{F}(z^0), (x^1, y^1) \in \mathcal{F}(z^1)$ respectively achieving bounds $g(y^0) \geq \gamma^0$ and $g(y^1) \geq \gamma^1$. For each term $u = x, y, z, \gamma$ just defined, we correspondingly define $u^q = (1 - q)u^0 + qu^1$, for $q \in [0, 1]$.

By concavity of g , convexity of f_j , $j = 1, \dots, m$, and linearity, we have

$$\begin{aligned} g(y^q) &\geq (1 - q)g(y^0) + qg(y^1) \geq \gamma^q, \\ f_j(x^q(\tau)) &\leq (1 - q)f_j(x^0(\tau)) + qf_j(x^1(\tau)) \leq c_j, \\ \mathbb{E}_\mu[x_i^q - y_i^q] &= (1 - q)z_i^0 + qz_i^1 = z_i^q, \end{aligned}$$

for $\tau \in \mathcal{T}, j = 1, \dots, m$ and $i = 1, \dots, n$. These above inequalities show that $(z^q, \gamma^q) \in \mathcal{C}$ and thus our claim is holds: \mathcal{C} is convex.

Let $\gamma^* = \max_{(x, y) \in \mathcal{F}} g(y)$. Here we are optimizing over $\mathcal{F}(z)$ with $z = 0$. So, it is clear that $(0, \gamma^*)$ does not belong to the interior of \mathcal{C} . Thus by the Supporting Hyperplane Theorem Rockafellar (1997), there exists a hyperplane through $(0, \gamma^*)$ supporting \mathcal{C} . In other words, there exists a non-zero vector $(b, \phi) \in \mathbb{R}^{\mathcal{I}} \times \mathbb{R}$ such that

$$\phi \gamma^* \geq \phi \gamma + b^\top z,$$

for all $(z, \gamma) \in \mathcal{C}$. Firstly, it is clear that $\phi \geq 0$, otherwise γ^* is not maximal for $(x, y) \in \mathcal{F}$.

We now claim $\phi \neq 0$. We proceed by contradiction. If $\phi = 0$, then $0 \geq b^\top z$ for all $(z, \gamma) \in \mathcal{C}$. But notice, for any $x \in \mathcal{B}(\mathcal{T}, \mathbb{R}_+^n)$, we can choose $y_i \in \mathbb{R}$ such that $y_i - \mathbb{E}_\mu[x_i] = b_i$, thus for this choice of (x, y) we have $z = b$. Thus, $b^\top z = b^\top b > 0$, and so we have a contradiction. It must be that $\phi > 0$.

As $\phi > 0$, we can define $b^* = (b_i/\phi : i \in \mathcal{I})$. Since for each $(x, y) \in \mathcal{B}(\mathcal{T}, \mathbb{R}^n) \times \mathbb{R}^{\mathcal{I}}$, if we set $z_i = \mathbb{E}_\mu[x_i - y_i]$ and $\gamma = g(y)$ then we have $(z, \gamma) \in \mathcal{C}$. With this we have

$$\max_{(x', y') \in \mathcal{F}} g(y') = \gamma^* \geq \gamma + b^{*\top} z = g(y) + \sum_{i \in \mathcal{I}} b_i^* \mathbb{E}[x_i - y_i]$$

Thus, maximizing over $x \in \mathcal{B}(\mathcal{T}, \mathbb{R}^n)$ and $y \in \mathbb{R}^{\mathcal{I}}$, we have

$$\max_{(x, y) \in \mathcal{F}} g(y) \geq \max_{\substack{y \in \mathbb{R}^n \\ x \in \mathcal{B}(\mathcal{T}, \mathbb{R}^n)}} g(y) + \sum_{i \in \mathcal{I}} b_i^* \mathbb{E}[x_i - y_i]. \quad (\text{C.4})$$

Together (C.3) and (C.4) give the required equality (C.2). In addition, given (C.1) has a finite optimum, inequality (C.4) can only hold when $b^* \geq 0$.