REAL-TIME BIDDING FOR TIME CONSTRAINED IMPRESSION CONTRACTS IN FIRST AND SECOND PRICE AUCTIONS – THEORY AND ALGORITHMS

A PREPRINT

R. J. Kinnear
ryan@kinnear.ca

R. R. Mazumdar
mazum@uwaterloo.ca

P. Marbach
marbach@cs.utoronto.edu

November 17, 2021

ABSTRACT

We study the optimal behavior of a bidder in a real-time auction subject to the requirement that a specified collections of heterogeneous items be acquired within given time constraints. The problem facing this bidder is cast as a continuous time optimization problem which we show can, under certain weak assumptions, be reduced to a *convex* optimization problem. Focusing on the standard first and second price auction mechanisms, we first show, using convex duality, that the optimal (infinite dimensional) bidding policy can be represented by a single *finite* vector of so-called "pseudo-bids". Using this result we are able to show that, in contrast to the first price auction, the optimal solution in the second price case turns out to be a very simple piecewise constant function of time. Moreover, despite the fact that the optimal solution for the first price auction is genuinely dynamic, we show that there remains a close connection between the two cases and that, empirically, there is almost no difference between optimal behavior in either setting. Finally, we detail methods for implementing our bidding policies in practice with further numerical simulations illustrating the performance.

Keywords Real-time Bidding; Infinite-dimensional Optimization; Computational Advertising; Auction Theory; First Price Auction; Second Price Auction; Transportation Production Problem

1 Introduction

Since at least 2016 online advertising, which includes ads in cellphone apps, website banner ads, search engine keyword ads, etc. has been the dominant advertising segment [19]. One of the drivers of this dominance is not only the fact that so many people are now online, but that the technology enables more accurate targeting. That is, advertisers can accurately target their messaging to the groups or individuals to which the message is most relevant. This technology enables advertisers to raise awareness of their products, website or app operators (i.e., "publishers") to monetize their audiences, and internet users to access a wealth of gratis content.

There are many different mechanisms through which publishers can sell ad space. For example, sponsored search is a well known mechanism for search engines where certain keywords (e.g. "mortgage", "used car", etc.) trigger the display of relevant sponsored results. More traditional fixed contracts (e.g., for the display of fixed banner ads) between advertisers and publishers also remain common, particularly for businesses which need to be careful about their image and want to maintain complete control over what their brand is associated with. Finally, a significant portion of total advertising revenue is generated through a mechanism called real-time bidding (RTB), which is the mechanism at the focus of this paper.

RTB is facilitated by auction exchanges (e.g., Google AdX [28]) where the arrival of each user to a publisher website or app triggers a sealed-bid auction (the bid is hidden from other participants) wherein advertisers bid on the right to display their ad to the particular user. Generically, we will think of this as a single "item", and the known characteristics (e.g., age, interests, etc.) of the visitor as the item's "type". Advertisers decide upon which types they wish to target, and bid on these items when they arrive at the auction exchange.

In this paper we study a problem of interest to intermediaries or brokers known as Demand Side Platforms (DSPs), see e.g. [39]. Advertisers contract with these DSPs to bid on the RTB market on their behalf. Specifically, we consider DSPs that have entered into contracts which require a specified number of items of particular types be won within a given time deadline. We do not allow for the contract requirements to be subject to post hoc negotiation – once the contract is accepted the DSP is obligated to meet the requirements. This benefits the advertiser by allowing them to offload the risk associated with price changes in the market, and gives the DSPs the opportunity to profit from managing this risk on behalf of multiple advertisers. Thus, the objective of the DSP is to win the required items in RTB at the lowest possible cost (since the difference between the contract value and its cost of acquiring the ads is the profit). This is a natural setting to consider as there is a reasonable expectation of efficiency gains when the demands of a large number of advertisers are aggregated by intermediaries.

The dominant auction mechanisms in RTB are *first price* and *second price* auctions. In either case, the item is awarded to the highest bidder, but the difference lies in what they pay: in a first price auction, the winning bidder pays their actual bid, whereas in a second price auction, the winning bidder pays whatever the *second highest* bid was. In both auctions, the bids are sealed (other participants don't observe them, even after an item is allocated) so the only source of information about prices are from available historical data, and the knowledge gained from the bidder's direct participation. Standard overviews of auction theory are provided by [24, 29].

Related Work Real-time bidding is commonly studied from at least two main perspectives: the descriptive analysis of the market itself (e.g., [20, 2, 3]), where questions of game theoretic equilibria are at the forefront; and the consideration of normative problems of optimal bidding. Our perspective is normative. We solve an optimal bidding and allocation problem by providing algorithms which produce a plan over a given time horizon specifying the optimal average rate that items should be acquired, and to which contracts (or campaigns) they should be allocated. Similar allocation problems have been considered from the publisher's viewpoint [8, 33] and the advertiser's [48, 37]. However, aside from [37], no other papers, to our knowledge, have considered contract management aspects with hard constraints from the DSP perspective. Many works have studied the problem of short time horizon adaptation, for instance the works of [16, 47, 21, 23] have applied classical control theory methods (e.g., PID controllers) and those of [18, 7, 43] MDP formulations and reinforcement learning. The goal of these works are to maintain or maximize certain performance indicators, e.g., click through rates or purchasing decisions. As such, these works can be seen as considering a faster time scale in comparison to our problem. Optimal bidding problems over a long time horizon are considered by [13, 17] via application of stochastic optimal control, where the goal is to maximize the total valuation of items attained over the period, subject to budget constraints. An important consideration has thus also been the estimation of item valuations, and how to map valuations into bids, see e.g., [32, 49]. These are distinct from the present work since we do not appeal to any notion of item valuation, indeed, a valuation can only be derived from the contract requirements. Our formulation is also dual to many of these as we seek to minimize the cost of obtaining a given number of items, as opposed to maximizing the utility of items obtained within a given budget. A similar class of problems to our optimal acquisition problem is that of revenue management (see e.g., [40, 6]) wherein a merchant needs to set prices for their wares in order to optimally liquidate their inventory. Aside from the superficial distinction of sell-side versus buy-side, the contract management problem studied in the present paper involves substantially different constraints. Finally, many papers have also studied separate, but important, statistical estimation problems (e.g. [11, 44, 51, 52, 45, 15]) concerned with estimating the prevailing prices of arriving items.

Contributions This paper studies an optimal bidding and allocation problem similar to the earlier work of [37] for second price static auctions, but with the additional consideration of temporal dynamics and differing contract deadlines. The main contributions of this paper over those of [37] are threefold: Firstly, the recognition that the problem is in fact naturally *convex*, this leads to an analysis which is general enough to unify both the first and second price auctions, as well as to result in an insightful duality theory. Secondly, a clear development of practical and simple algorithms for implementing solutions to the optimal bidding and allocation problem are provided, along with detailed simulation evidence and practical methods for adaptation to a stochastic environment. And finally, we show that there is a close connection between the first and second price cases, both theoretically as well as through experimental evidence.

Outline The outline of the paper is as follows: We formally introduce the problem in Section 2. In Section 2.3, we show that it can be reformulated as an infinite dimensional convex optimization problem (albeit with some additional

assumptions in the first price case, see Proposition 2.2). The convex formulation and a strong duality theorem is provided in Section 2.4, where we also show that the optimal time dependent bids are fully determined by a finite set of parameters deemed *pseudo bids*. Using duality, we show in Section 2.5 that there is a close connection between the first and second price case; this connection is expanded upon in the simulations of Section 3.2, where experimental evidence demonstrates that there is little practical difference between the solutions calculated in the two cases. In terms of implementation, the dual problem in Section 2.4 is finite and can be solved to attain exact solutions, but doing so requires computationally expensive integration steps; this motivates a simpler polyhedral approximation method in Section 3.1, where we also bound the error between the approximation and the true optimal solution. Using the methods of Section 3.1, an empirical study is carried out in Section 3.2 using real RTB data. Motivated by the results of this simulation, Section 3.3 examines additional practical methods for online bid updates and for accounting for uncertainty in any estimated quantities. Finally, Section 4 concludes. Proofs are provided in the appendix.

2 The General Optimal Bidding Problem

We consider the problem of bidding in a sequential auction with M available item types $j \in [M] \triangleq \{1, 2, \ldots, M\}$ and where we are required to fulfill N contracts by winning a specified number of items of given types. Specifically, the contract $i \in [N]$ is a 3-tuple $(\mathcal{A}_i, C_i, T_i)$, where $\mathcal{A}_i \subseteq [M]$ is the set of item types that can be allocated towards the contract, C_i is the number of items (of any type in $\mathcal{A}_i$) that are needed, and T_i is the time deadline by which the items must be obtained. Another collection of sets, $\mathcal{B}_i$, is uniquely determined by the condition $i \in \mathcal{A}_j \iff j \in \mathcal{B}_i$, and the sets $\mathcal{A}_i, \mathcal{B}_j$ therefore determine a bipartite graph $\mathcal{G}$ on $[N] \times [M]$ encoding which item types can be allocated to which contracts, and conversely, which contracts can be satisfied by which item types. It will be convenient to write $(i, j) \in \mathcal{G}$ to mean that i, j is an edge in this graph.

We will suppose without loss of generality that $0 < T_1 \leq T_2 \leq \cdots \leq T_N$. Contracts stipulating more complicated item requirements (e.g., at least 100 of type j, and an additional 50 of either j or j') can be implemented by introducing multiple contracts with the same deadlines. Denoting $T \triangleq T_N$ for the end of the problem horizon and $T_0 \triangleq 0$ will also be notationally convenient. The time T_i indicates when contract i "exits" the problem, and similarly it will be convenient to write $T^j = \max_{i \in \mathcal{B}_j} T_i$ for the time that items of type j exit, i.e., are no longer needed.

Remark 2.1 (Non-zero Initial Times). The framework can also incorporate introducing new contracts part way through the fulfillment of existing ones by "resetting" the existing contracts by subtracting the current time and already obtained supply from their requirements, and then adding the new contracts to the collection before recalculating the optimal bids. This is further explored in Section 3.3.1 for the purpose of adapting to a stochastic environment. Managing contracts with starting times other than 0 is also not a major difficulty: each contract must now come along with a starting time S_i *and* it's deadline T_i, the ultimate result being that integrals over the interval $[0, T_i]$ would need to be taken over $[S_i, T_i]$, and the dual problem considered in Section 2.4 will have additional variables. In some applications (e.g., processor speed scaling [14]), there are algorithms which can take advantage of special structure in the arrivals of jobs (e.g., that the jobs be first-in first-out: $S_i \leq S_j \implies T_i \leq T_j$) to achieve performance improvements, but such structure does not impact the methods we develop here.

Remark 2.2 (Item types). That a contract may be fulfilled with any arrangement of items in $\mathcal{A}_i$ (e.g., C_i items of a single type $j \in \mathcal{A}_i$ is just as good as an even distribution over all items in $\mathcal{A}_i$) appears at first as an unusual modelling artifact. However, it need not be the case that the counter party can distinguish between item types. For example, consider two contracts: contract 1 targets anyone aged $30 - 40$, and contract 2 targets anyone aged $20 - 40$. This would result in the DSP recognizing two types of items: people aged $20 - 30$ and people aged $30 - 40$. Contract 2 can indeed be fulfilled just as well with either type, but contract 1 can be fulfilled only with the second type.

2.1 Supply Rate Curves and Cost Functions

The model of how items are acquired relies on a function $W_j(x, t)$, called a *supply curve*, associated to each item type j and quantifying, instantaneously at time t, the expected number of items that would be won by bidding x. This function derives from the assumption that our bidder is a *price taker*, i.e., has no substantial market impact.

The function $W_j(x, t)$ should be thought of us an un-normalized (cumulative) distribution function with respect to the bid argument x, but as a rate with respect to the time argument t. This function needs to be estimated from available historical data (see e.g., [11, 44, 51, 41, 15] for work on this problem), and represents what is essentially a fluid approximation of the "landscape" of competing bids. It is important to emphasize that these curves represent *averages*, but that the environment is stochastic in actuality. Assuming that we only have access to the first order average statistics leads to a simpler problem, while at the same time also being a *weaker* assumption than if we were

to develop a full stochastic model that assumes further information about higher order statistics. Nevertheless, we will see in Section 3.3 some practical methods for adapting to the stochastic environment.

We remark that the purpose of including time-dependence is to enable the consideration of forecast changes in supply availability; in particular, there are natural cycles present in the average available supply throughout time: see e.g., [46], as well as Figure 1.

Additionally, we must introduce a *cost function* $f_j(x, t)$ which quantifies, instantaneously at time t, the expected cost of bidding x on every arriving item of type j. This function naturally depends on the auction mechanism, but also on the supply rate curve W_j, since the total amount spent depends upon the number of items which are actually won. In order to lighten the notation, we will often suppress the j index and t argument, writing $W(x)$, $f(x)$ etc. when it is clear from, or irrelevant to, the context.

For a first price auction, the cost function is the product of the bid and the supply rate curve, since you pay your bid:

$$f^{1st}(x) = xW(x)\mathbf{1}_{\mathbb{R}_+}(x), \tag{1}$$

where

$$\mathbf{1}_A(x) = \begin{cases} 1 & x \in A \\ 0 & \text{otherwise} \end{cases},$$

is the indicator for the set A.

In the second price case, the supply rate curve becomes an integrator, since we only pay the highest competing bid:

$$\begin{aligned} f^{2nd}(x) &= \int_0^x u\mathsf{d}W(u)\mathbf{1}_{\mathbb{R}_+}(x) \\ &= \left[xW(x) - \int_0^x W(u)\mathsf{d}u \right] \mathbf{1}_{\mathbb{R}_+}(x), \end{aligned} \tag{2}$$

where the first integral is to be understood as a Lebesgue-Stieltjes integral (see e.g. [36, ch. 6]) and we have integrated by parts. The inequality

$$f^{2nd} \leq f^{1st}, \tag{3}$$

is evident from the above since $f^{2nd}(x) = f^{1st}(x) - \left[\int_0^x W(u)\mathsf{d}u \right] \mathbf{1}_{\mathbb{R}_+}(x)$.

We will consider both first price and second price auctions as the two cases are amenable to a unified treatment. However, we often place our focus on the first price case, as its analysis requires an additional (albeit rather weak) assumption. We will occasionally omit the superscript $1st$ or $2nd$ (sometimes written as a subscript) depending on whether or not we want to emphasize a particular case. Moreover, it is important to recognize that the DSP must estimate the curve $W(x)$ in the auction they will actually participate in. Due to strategic differences between the first and second price cases, these functions will not be the same across auction types, but our analysis has no need to distinguish between them.

We make the following basic assumptions.

Assumption 2.1 (Supply Rate Curve Properties)**.** For every $t \in \mathbb{R}_+$, the functions $x \mapsto W_j(x, t)$ are strictly positive and continuous on $\mathbb{R}$, as well as being either strictly monotone increasing on an interval $(-\infty, \bar{x}_j] \subseteq \mathbb{R}$ and flat thereafter, or strictly monotone increasing on $\mathbb{R}$ and unbounded. Moreover, $\bar{s}_j(t) \triangleq \sup_x W_j(x, t)$ is either $+\infty$ or attained at some $\bar{x}_j(t) \in \mathbb{R}$. Finally, $W_j(x, t) \to 0$ as $x \to -\infty$.

The reasons for these assumptions, and whether or not they can be relaxed, is discussed in the following sequence of remarks.

Remark 2.3 (Attainment)**.** That the supremum of W must be obtained when finite is not a benign assumption. It technically excludes many models of the form $W(x, t) = \lambda(t)W(x)$ where $\lambda(t)$ is a supply rate and $W(x)$ is a cumulative distribution function with unbounded support (e.g., $1 - e^{-\gamma x}$). Of course, this is essentially only a theoretical restriction, and such models are reasonable in practice since one could choose some large but fixed maximum allowable bid $\bar{x}$ and truncate. If the supremum is not attained when finite, then the inverse function W^{-1} is not lower semi-continuous, which is an important property for rigorous proofs.
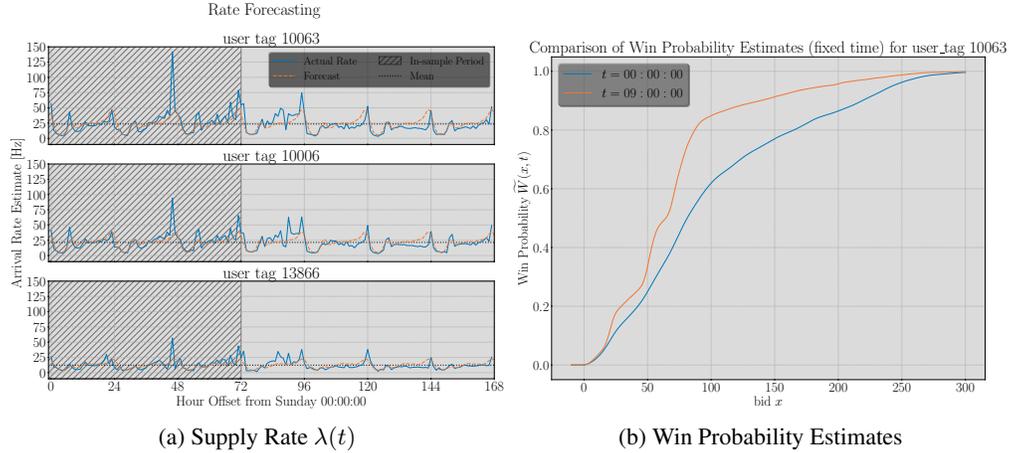
(a) Supply Rate $\lambda(t)$ (b) Win Probability Estimates

Figure 1: Estimated Supply Rate Curves and Costs

Supply rate and win probability curves estimated from iPinYou data. (a) Item arrival rates and the corresponding forecasts. The hatched region indicates an in-sample period with the remainder being out-of-sample. (b) Estimates of the win probability function $\widetilde{W}(x,t)$ are obtained via Gaussian Kernel Density Estimates.

Remark 2.4 (Randomized Bidding). It is known that there are benefits to *randomized bidding* [22]. That is, if $\sigma > 0$ is a parameter, and $x(t)$ is the estimated optimal bid at time t, then the *actual* bid submit to the auction house should be the random quantity $b \sim \mathcal{N}(x(t), \sigma^2)$. In this case, *nominal* bids below 0 still have some probability of inducing an *actual* bid above 0. The σ parameter can, for example, be chosen to correspond to the bandwidth of kernel density estimates of supply rate curves, a natural estimation procedure for this application. Thus, a nominal bid of $x \to -\infty$ can be understood as ensuring that the probability of the actual bid exceeding the floor price is converging to 0.

Remark 2.5 (Continuity). Following the discussion of the previous remark, randomized bidding also leads naturally to supply rate curves which are *smooth* (possessing higher order derivatives), for exactly the same reason that Gaussian kernel density estimates are smooth. Indeed, it would be reasonable to assume that supply rate curves are smooth, but we only *require* simple continuity. In the work of [37], supply curves are initially assumed merely to be only right-continuous; however, it is then shown that optimal bidders *must* randomize, which makes the curves effectively continuous.

2.1.1   Estimation of Supply Rate Curves

In order to provide further context, we briefly discuss estimates of supply curves obtained from the IPinYou dataset [50, 27]. We are not focused on the statistical estimation problem, and we therefore use simple methods. Further detail is provided in the appendix.

In this example, the supply curve $W(x,t)$ is estimated by combining estimates of the supply rate (the rate of items arriving), denoted by $\lambda(t)$, and the probability of winning an item at a certain price. We denote the probability of winning by $\widetilde{W}(x,t)$, and the supply rate curve estimate is then taken to be $W(x,t) = \lambda(t)\widetilde{W}(x,t)$. Figure 1a depicts representative examples of rate estimates and clearly depicts the typical daily cycles in supply rates. Figure 1b depicts examples of $\widetilde{W}(x,t)$ estimates. The curves are positive, continuous, monotone increasing, and truncated to a bid of $x = 300$, c.f., Assumption 2.1.

2.2   Problem Definition

Our goal is to find a *bid path* $x(t) \in \mathbb{R}^{N \times M}$ and *allocation path* $\gamma(t) \in [0,1]^{N \times M}$ which solves the following problem:

$$\underset{x,\gamma}{\text{minimize}} \quad \sum_{i=1}^{N} \int_0^{T_i} \Big[\sum_{j \in \mathcal{A}_i} \gamma_{ij}(t) f_j(x_{ij}(t), t) \Big] \mathrm{d}t$$

$$\text{subject to} \quad \sum_{j \in \mathcal{A}_i} \int_0^{T_i} \gamma_{ij}(t) W_j(x_{ij}(t), t) \mathrm{d}t \geq C_i \qquad (TP)$$

$$\sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) \leq 1, \gamma_{ij}(t) \geq 0,$$

where it is to be understood that constraints containing unqualified indices are to hold over their whole range. That is, $\gamma_{ij}(t) \geq 0$ is short for: $\forall (i,j) \in \mathcal{G} \; \forall t \in [0, T_i] \; \gamma_{ij}(t) \geq 0$. Additionally, functions $\gamma_{ij}(t), x_{ij}(t)$ etc. are treated as 0 or undefined outside of $t \in [0, T_i]$. The notation

$$\mathcal{T}_t \triangleq \{ i \in [N] \mid t < T_i \} \qquad (4)$$

indicating the contracts active up to time t will be useful in keeping track of active contracts.

The variable $\gamma_{ij}(t)$, referred to as the *allocation*, parameterizes the proportion of items of type j which will be bid on on behalf of contract i. Similarly, $x_{ij}(t)$ indicates the bid that should be submit for an item of type j on behalf of contract i. Concretely, upon the arrival of an item of type j at time t, the DSP bids on the item with probability $\Gamma_j(t) = \sum_{i \in \mathcal{B}_j} \gamma_{ij}(t)$, and if they have chosen to bid they do so on contract i with probability $\gamma_{ij}(t)/\Gamma_j(t)$, and finally place the bid $x_{ij}(t)$ on the item. We have here allowed for the possibility of placing different bids depending upon which contract the item will be allocated towards, however, it will be shown in the sequel that this is never necessary and that Problem (TP) could have been written with bids $x_j(t)$ indexed only by j and t, rather than the triple i, j, t.

We set the stage for our solution of Problem (TP) with some final conventions.

Definition 2.1 (Conventions). We write $W_j^{-1}(x, t)$ for the inverse function of $x \mapsto W_j(x, t)$. The inverse is guaranteed to exist on the range of W_j since the function is strictly monotone. Outside the range, the inverse is understood to be $+\infty$, i.e., $W_j^{-1}(s, t) = \infty$ if there is no $x \in \mathbb{R}$ such that $W_j(x, t) = s$. Typically, we will use x for bids and s for units of supply.

2.3 Convex Acquisition Costs

The key to understanding Problem (TP) is to consider *acquisition cost* functions:

$$\Lambda_j(s, t) \triangleq f_j \circ W_j^{-1}(s, t) \qquad (5)$$

which quantify, instantaneously at time t, the average cost of acquiring supply at the rate s.

In the second price case, the acquisition cost function is *always* convex. Indeed, this can be seen by a change of variables $y = W(u)$

$$\Lambda^{2nd}(s) = \mathbf{1}_{\mathbb{R}_+}(W^{-1}(s)) \int_0^{W^{-1}(s)} u \, \mathrm{d}W(u)$$

$$= \mathbf{1}_{\{s \geq W(0)\}} \int_{W(0)}^{s} W^{-1}(y) \, \mathrm{d}y,$$

which is convex since its (one-sided) derivative on the domain is monotone non-decreasing. The interpretation of this result is that the marginal acquisition costs are increasing [26]. This is entirely natural for second price auctions since you pay the same amount for items that would be won by bids of x, whether or not you're bidding exactly this amount, or something greater.

For the case of first price auctions, while total costs are necessarily always greater for greater bids, the *marginal* costs need not necessarily be increasing and therefore convexity does not always hold. In order to understand when it does, we need to discuss the idea of log-concavity. Recall that a function W is log-concave if $\ln \circ W$ is concave and that any concave function is log-concave, but that the converse is false. The following notion of concavity, which includes both concavity and log-concavity, is utilized in our analysis.

Definition 2.2 (α-concavity). Define, for $\alpha \geq 0$, $x > 0$ the function

$$\ell_\alpha(x) \triangleq \int_1^x \frac{1}{t^\alpha} \mathrm{d}t = \begin{cases} \ln x & \alpha = 1 \\ \frac{x^{1-\alpha}-1}{1-\alpha} & \text{otherwise} \end{cases},$$

where in particular $\ell_2(x) = 1 - 1/x$. We will say that a positive function $W : \mathbb{R} \to (0, \infty)$ is (strictly) α-concave if $\ell_\alpha \circ W$ is (strictly) concave. In particular, W is log-concave if $\alpha = 1$ and concave if $\alpha = 0$.

The following observation tells us that for any $\alpha > 1$, α-concavity is a weaker requirement than log-concavity, which is in turn weaker than concavity.

Proposition 2.1 (Hierarchy of α-concavity). *For $0 \leq \alpha < \beta$, if W is α-concave, then it is also β-concave.*

We can now characterize the conditions under which the cost of acquisition function $\Lambda_j(s, t)$ is convex in s for first price auctions. As discussed above, a similar statement holds for the second price case, and doesn't require the log-concavity assumption.

Proposition 2.2 (Convex Acquisition Costs – First Price Case). *Suppose that a supply rate curve $W(x)$ is strictly 2-concave. Then in a first price auction, the extended acquisition function*

$$\Lambda^{1st}(s) \triangleq \begin{cases} \infty; & s > \bar{s} \\ 0; & s < W(0) \\ sW^{-1}(s); & \text{otherwise} \end{cases}, \tag{6}$$

is a proper[1], lower semi-continuous, non-decreasing, and convex function on $\mathbb{R}$, which is strictly convex on $[W(0), \bar{s}]$. Moreover, the convex conjugate

$$\Lambda^*_{1st}(p) \triangleq \sup_{s \in \mathbb{R}} \big[sp - \Lambda(s)\big],$$

is itself a proper, lower semi-continuous, non-decreasing, and convex function on $\mathbb{R}$, which satisfies $\Lambda^*_{1st}(p) = \infty$ for $p < 0$, and if $\bar{s} < \infty$, then $\Lambda^*_{1st}(p) = p\bar{x} - \Lambda_{1st}(\bar{s})$ on $p > \bar{x}$.

2.4 Optimal Bidding in Continuous Time

As long as the acquisition cost functions are convex, we can equivalently reformulate Problem (TP) as the following continuous time *convex* optimization problem. This reformulation is not trivial and makes use of the existence of a campaign-independent optimal bid $x_{ij}(t) = x_j(t)$, the existence of which is established by using the convexity of Λ.

Proposition 2.3 (Continuous Primal Problem). *In a first or second price auction, suppose that for each $j \in [M]$ and $t \in [0, T^j]$, the acquisition cost curve $\Lambda_j(x, t)$ is convex in x. Then, Problem (TP) can be reformulated as*

$$\begin{aligned}
\underset{s,r}{\text{minimize}} \quad & \sum_{j=1}^M \int_0^{T^j} \Lambda_j(s_j(t), t) \mathrm{d}t \\
\text{subject to} \quad & \sum_{j \in \mathcal{A}_i} \int_0^{T_i} r_{ij}(t) \mathrm{d}t \geq C_i \\
& \sum_{i \in \mathcal{B}_j} r_{ij}(t) = s_j(t) \\
& r_{ij}(t) \geq 0.
\end{aligned} \tag{P}$$

If a solution exists, then a solution to the original problem (TP) is obtained via $x_{ij}(t) = W_j^{-1}(s_j(t), t)$ for each $i \in \mathcal{B}_j$ and $\gamma_{ij}(t) = r_{ij}(t)/s_j(t)$. Moreover, Problem (P) is convex.

Remark 2.6 (Connection to Transportation-Production Problems). This result demonstrates that Problem (TP) can be reformulated as an instance of a (convex) *Transportation-Production problem* (see [26, 35]) in continuous time

---

[1]Recall that a function $f : \mathbb{R} \to (-\infty, \infty]$ is *proper* if it is not everywhere equal to $+\infty$.

(hence the label (TP)), which is in turn an example of a *monotropic program* [34, 4]. To see this connection, one can think of i as the "consumption nodes", j as the "production nodes", the graph $\mathcal{G}$ (defined through the sets $\mathcal{A}_i, \mathcal{B}_j$) as defining the transportation network, and the functions f_j, W_j, Λ_j as characterizing the costs of production.

In order for a solution to exist, it is necessary only that an adequate amount of supply is available to fulfill each contract before its deadline. In the context of internet advertising, this is not unreasonable, and a coarse sufficient condition is $W(x) \to \infty$ as $x \to \infty$. As in the case of finite optimization problems, strong duality obtains if there is a *strictly* feasible point.

Assumption 2.2 (Adequate Supply). We suppose that $\exists x \in \mathbb{R}^M$ such that the following linear feasibility problem has a strictly feasible point (i.e., a Slater point)

$$
\begin{aligned}
\text{find} \quad & \gamma_{ijk} \\
\text{such that} \quad & \sum_{k:T_k \leq T_i} \sum_{j \in \mathcal{A}_i} \gamma_{ijk} \int_{T_{k-1}}^{T_k} W_j(x_j, t)\mathrm{d}t > C_i \\
& \sum_{i \in \mathcal{B}_j} \gamma_{ijk} \leq 1, \gamma_{ijk} \geq 0.
\end{aligned}
\tag{7}
$$

That is, there exists a large enough constant bid such that an adequate amount of supply is attained.

Remark 2.7. A simpler sufficient condition (appearing also in [26]) is to suppose that each curve has an adequate amount of supply to fulfill *all* the contracts to which it is assigned:

$$
\forall j \in [M] \ \exists x \in \mathbb{R} : \int_0^{\tau_j} W_j(x, t)\mathrm{d}t > \sum_{i \in \mathcal{B}_j} C_i,
$$

where $\tau_j = \min\{T_i \mid i \in \mathcal{B}_j\}$. Finally, this condition is clearly satisfied under the simple condition that if for each j, t we have $W_j(x, t) \to \infty$ as $x \to \infty$.

Proposition 2.4 (Duality). *A Dual of (P) can be formulated as*

$$
\begin{aligned}
\underset{\rho, \mu}{\text{maximize}} \quad & -\sum_{j=1}^M \sum_{k:T_k \leq T^j} \int_{T_{k-1}}^{T_k} \Lambda_j^*(\mu_{jk}, t)\mathrm{d}t + \sum_{i=1}^N \rho_i C_i \\
\text{subject to} \quad & \mu_{jk} \geq \rho_i \ \forall i \in \mathcal{B}_j \cap \mathcal{T}_{T_k} \\
& \rho_i \geq 0,
\end{aligned}
\tag{D}
$$

which is a finite convex problem. Problem (D) is dual to Problem (P) in the sense that if $D^\star$ and $P^\star$ are their respective values (possibly ∞ or $-\infty$), then $D^\star \leq P^\star$.

Finally, under Assumption 2.2 there exists a solution $(s, r) \in L_2(\mathbb{R})^M \times L_2(\mathbb{R})^{N \times M}$ to Problem (P) and a solution $(\rho, \mu) \in \mathbb{R}^N \times \mathbb{R}^M$ to Problem (D) and $-\infty < D^\star = P^\star < \infty$.

Finally, the optimal bid path $x(t)$ for Problem (TP) can be reconstructed from solutions to (D) and is, in fact, fully characterized by the finite vector $\rho \in \mathbb{R}^N$, which we designate as the vector of *pseudo-bids*.

Example 2.1 (Pseudo Bids). Suppose that μ_{jk}, ρ_i are optimal dual multipliers for the Problem (D) and consider the s-dependent portion of the Lagrangian:

$$
L(s) = \sum_{j=1}^M \int_0^{T^j} \left[\Lambda_j(s_j(t), t) - \mu_j(t)s_j(t) \right] \mathrm{d}t,
$$

where $\mu_j(t) = \mu_{jk}$ for $t \in [T_{k-1}, T_k)$ is piecewise constant.

Minimizing this function over $s_j(t)$ leads to an optimal supply path as a function of $\mu_j(t)$. Since the above function is convex, we know that the optimum is attained at $s_j(t)$ such that the pointwise selection holds:

$$
\mu_j(t) \in \partial \Lambda_j(s_j(t), t),
\tag{8}
$$

where ∂ indicates the subdifferential and is with respect to the first argument. Therefore, since $\mu_j(t) = \max\limits_{i \in \mathcal{B}_j \cap \mathcal{T}_t} \rho_i$, the entire optimal supply path $s(t)$, and therefore the entire optimal bid path $x(t)$, is fully determined by the *finite* vector $\rho \in \mathbb{R}^N$ through Equation (8). For this reason, we refer to ρ as the vector of *pseudo bids*.

2.5   Connections Between First and Second Price Auctions

It will be seen in Section 3.2 that, empirically, the optimal solution calculated for the second price auction also provides very good bid and allocation paths for the first price auction. From Example 2.1, we know that the optimal solutions cannot be exactly the same (except perhaps in some coincidental cases), but we explain the empirical similarity by illuminating the close connection between the first and second price cases in the following proposition.

Proposition 2.5 (Connection Between First and Second Price Case). *Under Assumption 2.2, and whenever the acquisition cost curves are convex, the Problem (TP) for a second price auction is equivalent to solving*

$$\begin{aligned}
\operatorname*{minimize}_{x,\gamma,\rho} \quad & \sum_{j=1}^{M} \int_{0}^{T^j} x_j(t) W_j(x_j(t), t)\mathrm{d}t - \sum_{i=1}^{N} \rho_i C_i \\
\text{subject to} \quad & \sum_{j \in \mathcal{A}_i} \int_{0}^{T_i} \gamma_{ij}(t) W_j(x_j(t), t) \geq C_i \\
& \sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) = 1, \gamma_{ij}(t) \geq 0, \rho_i \geq 0 \\
& x_j(t) \geq \rho_i \ \forall i \in \mathcal{B}_j \cap \mathcal{T}_t.
\end{aligned} \tag{9}$$

The objective function involves $x_j(t) W_j(x_j(t), t)$, which is the cost associated with a first price auction.

Returning to Example 2.1, we can make a further explicit comparison between the two cases.

Example 2.2 (Pseudo Bids (continued)). For the purpose of illustration, let us temporarily assume that our supply rate curves are smooth in the bid x, and that the optimal solutions satisfy $s_j(t) \in (W_j(0,t), \bar{s}_j(t))$ (i.e., do not reach the extremes of available supply) so that we have the simple expressions $\Lambda_j^{2nd}(s,t) = \int_{W_j(0,t)}^{s} W_j^{-1}(y,t)\mathrm{d}y$ and $\Lambda_j^{1st}(s) = s W_j^{-1}(s,t)$. For the first price case, let us define a function g_j, where $'$ denotes differentiation with respect to x:

$$g_j(x,t) \triangleq \frac{f'_{j,1st}(x,t)}{W'_j(x,t)} = x + \frac{W_j(x,t)}{W'_j(x,t)}.$$

We can then solve Equation 8 for the optimal bid in terms of μ_j, which in the first price case, can be written in terms of g: $\Lambda_j'^{-1}(x,t) = W_j \circ g_j^{-1}(x,t)$. The second price case is much simpler, and we have $\Lambda_{2nd}'(x,t) = W^{-1}(s,t)$. That is, the analogous 'g' function for the second price case is the identity.

Now, in each case we map the desired supply s into the required bid x through $x_j(t) = W_j^{-1}(s_j(t), t)$ and find that the optimal bids for each case are

$$x_j(t) = \begin{cases} \mu_j^{2nd}(t) & \text{Second Price} \\ g_j^{-1}(\mu_j^{1st}(t), t) & \text{First Price,} \end{cases} \tag{10}$$

where $\mu_j^{2nd}(t)$ is the dual multipliers associated to Problem (D) with second price cost function and μ_j^{1st} are multipliers of the same problem but with the first price cost function; that is, $\mu_j^{1st} \neq \mu_j^{2nd}$ in general. It is important to recognize that $\mu(t)$ is a piecewise constant function which is fully determined by $\rho \in \mathbb{R}^N$. The implication being that in each case, the entire time-dependent optimal bid path is fully characterized by the finite vector ρ associated to the dual problem (D).

The upshot of these calculations is that in the second price case, the optimal bid path is constant between contract deadlines, whereas in the first price case the pseudo bids need to be mapped through the function g^{-1}. The intuition for this fact is that in a second price case you are "protected" from overpaying (since you only pay the second highest bid) whereas in the first price case you need to adapt the bids to your belief about the prevailing prices.

If the supply rate curves are time-independent (i.e. $W(x,t) = W(x)$), then the function g_j^{-1} is time-independent as well. It may be conjectured that in this case the optimal bids would coincide, i.e., that $\mu_j^{2nd}(t) = g_j^{-1}(\mu_j^{1st}(t))$. Experimental evidence shows that this does not seem to be the case in general.

3   Computational Methods and Simulations

In this section we review computational methods useful for the implementation of our theory (Section 3.1); carry out a Monte Carlo simulation lending evidence to the similarity between the first and second price case, as suggested by Proposition 2.5 (Section 3.2); and finally provide basic results on practical methods for adapting to a stochastic environment (Section 3.3).

3.1   Computing and Implementing Optimal Bids

To summarize the developments so far, the solution to the optimal bidding and allocation problem (TP) consists of a bid path $x(t)$ and allocation path $\gamma(t)$. The function $\gamma_{ij}(t)$ indicates that if an item of type j is won at auction at time t, then it should be allocated towards contract i with probability $\gamma_{ij}(t)$. Therefore, upon the arrival of an item of type j, the bidder should sample a random index $\hat{i} \in \mathcal{B}_j$ from a categorical distribution with $\mathbb{P}\{\hat{i} = i\} = \gamma_{ij}(t)$. Then, the function $x_{\hat{i}j}(t)$ indicates the nominal bid which should be submit to the auction house (the actual bid being $\mathcal{N}(x_{\hat{i}j}(t), \sigma^2)$ if randomization is used). Finally, if the item is won, it should be allocated to contract $\hat{i}$.

The Problem (TP) is, as written, not tractable. However, Proposition 2.3 tells us that the solution $s(t), r(t)$ to the *convex* problem (P) can be used to obtain the optimal bid and allocation through

$$x_{ij}(t) = W_j^{-1}(s_j(t), t),$$
$$\gamma_{ij}(t) = r_{ij}(t)/s_j(t),$$

where it is to be understood that $x_{ij}(t) = \gamma_{ij}(t) = 0$ for any $j \notin \mathcal{A}_i$ or $t > T_i$, and that any $0/0$ encountered in the construction of $\gamma_{ij}(t)$ is to be taken as 0. Moreover, this formula tells us that we actually need to keep track of only M bidding functions, $x_j(t)$, since there is an optimal solution such that $x_{ij}(t) = x_{i'j}(t)$ for each $i, i' \in \mathcal{B}_j$.

Problem (P) is still an infinite dimensional problem, and thus remains intractable as written. A solution could be obtained by first solving the dual (D) (which is finite, but requires integration), substituting the resulting pseudo-bids (through g^{-1}, see Example 2.2) into the original problem (TP) and solving the resulting linear program to obtain $\gamma(t)$. However, we have found that directly discretizing the primal problem and solving the resulting convex program tends to be faster and with limited degradation in accuracy. Moreover, solving the primal directly results in *both* the optimal bids and the optimal allocations simultaneously.

Therefore, in this section we study the discretization of Problem (P). We will see that the linearity of the constraints allows us to construct *feasible* approximate solutions for (P), and that we can also estimate the sub-optimality of the approximation. Some additional details pertaining to the convexity of Λ and how the supply curves can be guaranteed to have the needed α-concave property are provided in the Appendix A. The conclusion of our developments is provided in Algorithm 1. The `error` condition in Algorithm 1 encountered when supply is inadequate can be handled, for example, by modifying the objective of Problem (P_K) to instead penalize supply shortfalls, rather than attempting to enforce them as constraints, resulting in a "best effort" solution.

To the end of discretizing Problem (P), choose a sequence of points $0 = \widetilde{T}_0 < \widetilde{T}_1 < \cdots < \widetilde{T}_K = T$ where $K \geq N$ and $\{T_i\}_{i=1}^N \subseteq \{\widetilde{T}_k\}_{k=1}^K$. Then, using the notation $\Delta_k = \widetilde{T}_k - \widetilde{T}_{k-1}$ we define, for $k \in [K]$ and $\widetilde{T}_k \leq T^j$

$$\overline{\Lambda}_{jk}(s) \triangleq \frac{1}{2}\Delta_k\big[\Lambda_j(s, \widetilde{T}_k) + \Lambda_j(s, \widetilde{T}_{k-1})\big], \tag{11}$$

i.e., a trapezoidal approximation of the integral $\int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} \Lambda_j(s, t)\mathrm{d}t$. The following finite approximation of Problem P is natural:

Proposition 3.1 (Finite Primal Problem). *The finite optimization problem over the variables $r_{ij}[k], s_{ij}[k]$ defined by*

$$\underset{s,r}{\text{minimize}} \quad \sum_{j=1}^{M} \sum_{k:\widetilde{T}_k \leq T^j} \overline{\Lambda}_{jk}(s_j[k])$$

$$\text{subject to} \quad \sum_{j \in \mathcal{A}_i} \sum_{k:\widetilde{T}_k \leq T_i} \Delta_k r_{ij}[k] \geq C_i \qquad\qquad (P_K)$$

$$\sum_{i \in \mathcal{B}_j \cap \mathcal{T}_{T_k}} r_{ij}[k] = s_j[k]$$

$$r_{ij}[k] \geq 0,$$

is a finite approximation of (P): any solution $(s_j[k], r_{ij}[k])$ of (P_K) is feasible for Problem (P) (as a piecewise constant function of t).

Let $g_j^{-1}(\mu, t)$ denote a continuous selection from Equation (8) of Example 2.1 (e.g., choose the right derivative). If each cost function $f_j(x,t) = xW_j(x,t)$ is L_f-Lipschitz in x (uniformly in t a.e.) and $g_j^{-1}(x,t)$ is Γ_g-Lipschitz in t a.e. (uniformly in x), then the cost difference ϵ_K of the optimal solution to (P) and the feasible piecewise constant approximation obtained from Problem (P_K) is bounded by

$$\epsilon_K \leq \frac{M}{4} L_f \Gamma_g \sum_{k=1}^{K} \Delta_k^2.$$

If for each j the function $W_j(x,t)$ (and hence $\Lambda_j(s,t)$) is almost everywhere twice continuously differentiable in t, the integral approximation error δ_K between the objective of (P) and (P_K) is bounded by $\delta_K = \mathcal{O}(\sum_{k=1}^{K} \Delta_k^3)$.

In particular, if we have $\Delta_k = T/K$ (valid as long as $\widetilde{T}_k$ contains the contract deadlines), then we have the bounds

$$\epsilon_K \leq \frac{M L_f \Gamma_g T^2}{4K}$$

$$\delta_K = \mathcal{O}(1/K^2).$$

The total discretization error between the value of (P_K) and the value of (P) is therefore $\epsilon_K + \delta_K = \mathcal{O}(1/K)$.

A summary of our proposed bidding methods are provided in Algorithm 1. The cost functions must correspond to either a second or a first price auction, and the supply curves must be strictly $2-$concave in the latter case, see Proposition 2.2.

Algorithm 1: Computing Optimal Bids

---

input : Contracts $\{(\mathcal{A}_i, C_i, T_i)\}_{i=1}^{N}$, supply curves $\{W_j(x,t)\}_{j=1}^{M}$, cost functions $f_j(x,t)$, parameter $K \geq N+1$

output : Bid path $x(t)$ and allocation path $\gamma(t)$

1 $(\widetilde{T}_k)_{k=0}^{K-N} \leftarrow \texttt{segment}([0,T], K-N)$ // cut $[0,T]$ into $K-N$ equal segments
2 $(\widetilde{T}_k)_{k=0}^{K} \leftarrow \texttt{sort}(\{\widetilde{T}_k\}_{k=0}^{K-N} \cup \{T_i\}_{i=1}^{N})$ // incorporate contract deadlines
3 Let $\Lambda_j(s,t) = f_j \circ W_j^{-1}(s,t)$ // acquisition function
4 **if** *Assumption 2.2 holds (adequate supply is available)* **then**
5 $\quad \lfloor \; s_j[k], r_{ij}[k] \leftarrow \texttt{solve}(P_K)$ // solve discretized problem
6 **else**
7 $\quad \lfloor \;$ **return** error

8 $x_j(t) \leftarrow W_j^{-1}(s_j[k], t) \; \forall t \in [\widetilde{T}_{k-1}, \widetilde{T}_k), k : \widetilde{T}_k \leq T^j, j \in [M]$ // construct bid path
9 $\gamma_{ij}(t) \leftarrow r_{ij}[k]/s_j[k] \; \forall t \in [\widetilde{T}_{k-1}, \widetilde{T}_k), k : \widetilde{T}_k \leq T^j, i \in \mathcal{B}_j, j \in [M]$
10 **return** $(x(t), \gamma(t))$

3.2 Monte Carlo Simulations

We have run experiments to empirically evaluate the performance of our methods on real data from the well known
IPinYou dataset [27, 50]. Further details on the specific simulation methodology are provided in the appendix. Our
simulations in this section focus on the similarity between the optimal solutions computed in the first and second price
settings cf. Proposition 2.5.

All our computations have been carried out with Python's scientific computing stack [38] and CVXPY [12] [1]. To
summarize our setup, we have estimated supply rate curves $W_j(x, t)$ via Gaussian kernel density estimation using the
one week of available data stratified into twenty-four one hour intervals. The bandwidth is chosen via Silverman's rule
(see e.g., [42, ch. 6]), which results in smooth estimates, and is not likely to be overfit. The hourly stratification results
in a supply rate curve estimate which accounts for daily periodic trends, but not weekly ones. In practice, supply curve
estimation is made more challenging by the fact that price observations are *censored*, that is, the true winning price is
only observed by the winner [44, 51, 45]. However, the IPinYou dataset was obtained by submitting very high bids
which win almost every impression, and we are therefore able to comfortably ignore the affects of censoring in our
simulation with the understanding that more sophisticated supply rate curve estimation methods are to be applied in
practice.

The problems we simulate in this section contain $N = 6$ contracts and $M = 5$ item types. Each contract in the
collection is randomly sampled uniformly at random within prescribed bounds. In particular, the starting "real" time
point (i.e., time $T_0 = 0$) is sampled from anywhere between the bounds of available data, the length of contracts
is uniformly random between 0 and 70 hours, and the number of required items are sampled uniformly between a
small (easily fulfilled) lower bound, and a large (near the maximum available supply) upper bound – random contracts
which are not feasible are re-sampled. A total of 250 Monte Carlo iterations are carried out. We use the same estimated
supply rate curves in each simulation, but the auction prices and item arrival times are sampled directly from real data[2].

Whenever an adequate number of items to fulfill a contract are acquired, that contract is removed from simulation
and a completely new set of bids are calculated given the new (reduced) set of contracts. Not doing so would result
in *overfilling* some contracts, which would never be done in practice. Thus, since there are a total of 6 contracts, we
solve 6 instances of Problem (P_K) over the course of each simulation run.

When the problem data (i.e., the supply rate curves W_j) have been estimated from the IPinYou dataset, there is almost
no difference between the results obtained for first or second price cases. That is, the solution (x, γ) for Problem
(TP) when $f = f_{2nd}$ is very nearly optimal for the same problem when $f = f_{1st}$, and vice versa. This similarity
is explained by the close connection between the two problems established by Proposition 2.5 – essentially, the main
difficulty is in finding bids which result in a feasible allocation. Note however, that all else being equal, second price
auctions are naturally much cheaper than first price auctions; in our case, the total spend in a first price auction is
on average 60% greater than in a second price auction. However, this obviously does not account for the incentive
differences of participants in different auction settings.

Synthetic examples where the solutions between first and second price auctions differ significantly can be constructed
and revolve around the function $g(x, t)$, identified in Example 2.1, which is used in the mapping from the constant
pseudo bids (dual multipliers) into (actual) bids $x(t)$. Since in the second price case $g_{2nd}^{-1}(p, t) = p$ for every t, the
most striking differences occur when g fluctuates rapidly with time. This corresponds to the case where the marginal
cost of production $f'_{1st}(x, t)/W'(x, t)$ (if the curves are differentiable in x) vary rapidly, as can be seen in Example
2.1. Major fluctuations of this sort are not observed in our dataset. We also note that examples where the solutions
differ, but the supply rate curves are time homogeneous (i.e., $W_j(x, t) = W_j(x)$) can also be constructed.

A summary of these numerical results is provided in Table 1. The simulation is carried out where the *true* auction
mechanism is a first price auction. The first and second row in the table corresponds to a cost metric (relative cost per
item acquired) with the third row measuring the mean fulfillment. The first row measures the average spend per item
over simulations that *actually* fulfilled all contracts (defined as reaching at least 98% fulfillment) and the second row
measures the average spend per item over all simulations. The first two columns in the table correspond to the two
possible cost functions used in solving Problem (P_K) (either a first or second price cost) and the third column provides
a p-value comparing the significance of the differences in the other two columns. The cost is measured relative to the
first column.

It would be expected that using a cost function which corresponds to the true auction mechanism would result in
improved performance. However, no real difference is observed. These results indicate that in practice, regardless
of whether the true auction mechanism is first price or second price, the bidder can bid *as if* it were a second price

---

[2]All of the code used in our experiments will be made available on github @RJTK.

Performance Metric	(TP) Cost Function		t-test p
	$f = f_{1st}$	$f = f_{2nd}$	
Spend / Item (fulfilled)	1	0.996	0.952
Spend / Item (total)	1	1.02	10^{-6}
Mean Fulfillment %	0.895	0.895	0.777

Table 1: Summary of Simulation Results (N = 250)

Relative average spend and mean normalized contract fulfillment in a monte carlo simulation of a first price auction (rows) for algorithms based on first price and second price cost functions (columns). (total) indicates averages taken over the entire set of simulations, and (fulfilled) indicates averages over fulfilled contracts only. The performance between the two algorithms is nearly identical. The p-values are obtained from paired sample t-tests `scipy.stats.ttest_rel` and independent sample t-tests `scipy.stats.ttest_ind` in the case of the first row (since there is an unequal number of samples in that case).

auction. This is advantageous not only for simplicity, but because Problem (P) is *always* convex for $\Lambda = \Lambda_{2nd}$, but not necessarily for Λ_{1st}, see Section 2.3.

The fulfillment row of Table 1 indicates the mean (over the 250 simulations) of the average (over the 6 contracts) percentage fulfillment. That is, if we let $c_i^{(k)}(t)$ be the cumulative number of items acquired for contract i by time t in the k^{th} simulation, we are reporting the mean over k of

$$C_{\text{avg}}^{(k)} \triangleq \frac{1}{N} \sum_{i=1}^{N} \frac{c_i^{(k)}(T_i)}{C_i}. \tag{12}$$

We turn to a further analysis of fulfillment in the next section.

3.3 Improving Fulfillment

Table 1 indicates that, on average, contracts are only about 90% fulfilled by their deadline. This is an important issue, but is not surprising because supply curves are averages while actual items arrive randomly. The theory that has been developed is thus optimal for the average case. Directly addressing the stochastic nature of the problem, e.g., via a dynamic stochastic model, is one possible alternative. However, such models are generally intractable, and do not easily lead to practical bidding algorithms. On the other hand, while staying within an average case framework, there are at least two effective methods for improving upon fulfillment: receding horizon updates, and risk-aware supply adjustments, which we consider in turn.

3.3.1 Receding Horizon

Suppose that at time t there has been $c_i(t)$ supply acquired for contract i. This is useful information which can be incorporated back into updated supply curve estimates, as well as into the calculation of future bids and allocations. Since we are assuming the use of only crude supply curve estimates from historical data, we consider only how to incorporate this information into new bid calculations. Most straightforwardly, a new instance of Problem TP can be solved for the updated contracts $\{(\mathcal{A}_i, C_i - c_i(t), T_i - t)\}_{i=1}^{N}$, and this bid re-calculation carried out at regular time intervals. This method of adaptation is referred to as a *receding horizon* method [25] and accounts for unexpected supply shortfalls by increasing the bid (which gets us closer to complete fulfillment), and decreases the bid when unexpected supply surpluses are encountered (which reduces costs).

3.3.2 Supply Inflation

In addition to a receding horizon, it is important to account for uncertainty in supply curve estimates, as well as the random dynamics intrinsic to the market. Our main algorithm does not endogenously account for the possibility of supply shortfalls, (it pertains only to an *average* case scenario) but a parameter, δ, can be appropriately tuned to inflate the supply.

Concretely, if C is the required supply, we can instead *aim* at an inflated supply target by solving Problem (P_K) with an inflated supply constraint $(1 + \delta)C$. The parameter δ can be tuned via cross-validation procedures, or calculated to correspond with a model of estimation and market uncertainty. This adjustment has the effect of front-loading the

acquisition of supply, a result similar to that of [17]. In combination with a receding horizon, excess spending will also be limited.

In order to justify this adjustment, we provide a simple example model. Suppose that $M = N = 1$ (i.e., there is a single item type and a single contract), that we have the contract $(\{1\}, C, T)$, and that the supply curve is now the *random and time-independent* quantity $\mathbf{W}(x)$ with $\mathbb{E}\mathbf{W}(x) = W(x)$, for some known $W(x)$. Then, given a risk tolerance parameter $\epsilon \in (0, 1)$ we seek a solution to the following problem

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & Tf(x) \\ \text{subject to} \quad & \mathbb{P}\big(T\mathbf{W}(x) \geq C\big) \geq 1 - \epsilon, \end{aligned} \tag{13}$$

where $f(x)$ is the expected cost per unit time in the auction when bidding x. Since the objective function is monotone increasing, the optimal solution to the problem is the smallest x which is feasible and we have established the following proposition.

Proposition 3.2. *Suppose that $f(x)$ is monotone non-decreasing. Then, the optimal solution of Problem* (13) *is given by*

$$x^\star(\epsilon) = \min\{x \mid \mathbb{P}\big(\mathbf{W}(x) \geq C_i/T\big) \geq 1 - \epsilon\}. \tag{14}$$

In principle, this result covers every probabilistic model for the distribution of $\mathbf{W}$ including the accounting of uncertainties in prices, item arrival rates, and estimation inaccuracies. Moreover, many such models for the function $x \mapsto \mathbb{P}\big(\mathbf{W}(x) \geq C_i/T\big)$ can be expected to be monotone non-decreasing in x (similarly to f), in which case the actual computation of $x^\star(\epsilon)$ poses no difficulty. It will be seen in the following example that plausible probabilistic models can reduce into the framework of supply inflation, i.e., aiming for some $(1 + \delta)C$.

Example 3.1 (Poisson Model). Suppose that $\mathbf{W}(x) \sim \text{Po}(W(x))$ has a Poisson distribution with mean $W(x)$. Although proposition 3.2 still applies to this model, and a computational solution is easy to obtain, we will make use of the Poisson Chernoff bound to find an illustrative approximation $\tilde{x}(\epsilon)$ in closed form. Replacing the probabilistic constraint in Problem (13) with the right hand side of the bound

$$1 - \mathbb{P}\big(\mathbf{W}(x) \geq C/T\big) \geq 1 - \frac{(eW(x))^{C/T} e^{-W(x)}}{(C/T)^{C/T}},$$

which is valid as long as $W(x) \geq C/T$, results in a *more* stringent requirement. Thus, a feasible point $\tilde{x}(\epsilon)$ for the problem can be found by solving the following inequality:

$$\frac{(eW(x))^{C/T} e^{-W(x)}}{(C/T)^{C/T}} \leq \epsilon$$

$$\overset{(a)}{\iff} x \geq W^{-1}\Big(-\frac{C}{T}\mathcal{W}_{-1}\big(-\frac{1}{e}\epsilon^{T/C}\big)\Big),$$

where in (a) $\mathcal{W}_{-1}$ is the lower branch of the Lambert-W function, i.e., the inverse function of $x \mapsto xe^x$, which is decreasing and negative [10]. Since $-\mathcal{W}_{-1}\big(-\frac{1}{e}\epsilon^{T/C}\big) \geq 1$, this has the form of a *multiplicative* inflation of the required supply, and can be written as $(1 + \delta)C$ where $\delta = -\mathcal{W}_{-1}\big(-\frac{1}{e}\epsilon^{T/C}\big) - 1$.

3.3.3   Sensitivity Simulations

In order to illustrate the effects of a δ risk adjustment and the receding horizon re-calculations, we have carried out simulations similar to those of Section 3.2, however with smaller contracts $N = M = 3$. The results are summarized by Figures 2 and 3.

Figure 2 provides a box plot comparing the parameter δ (applied to inflate the supply of each contract by the same proportion) to the average contract fulfillment (see $C_a^{(k)}$, Equation (12)). As is expected, the proportion of contracts which are completely fulfilled increases with δ. In addition, the plot provides pairs of boxes comparing the results when the bids are updated only after a contract is fulfilled (this is the minimal update schedule that doesn't overfill contracts), and when the bids are also updated every hour of simulation time (typically, about 20 updates are calculated over the
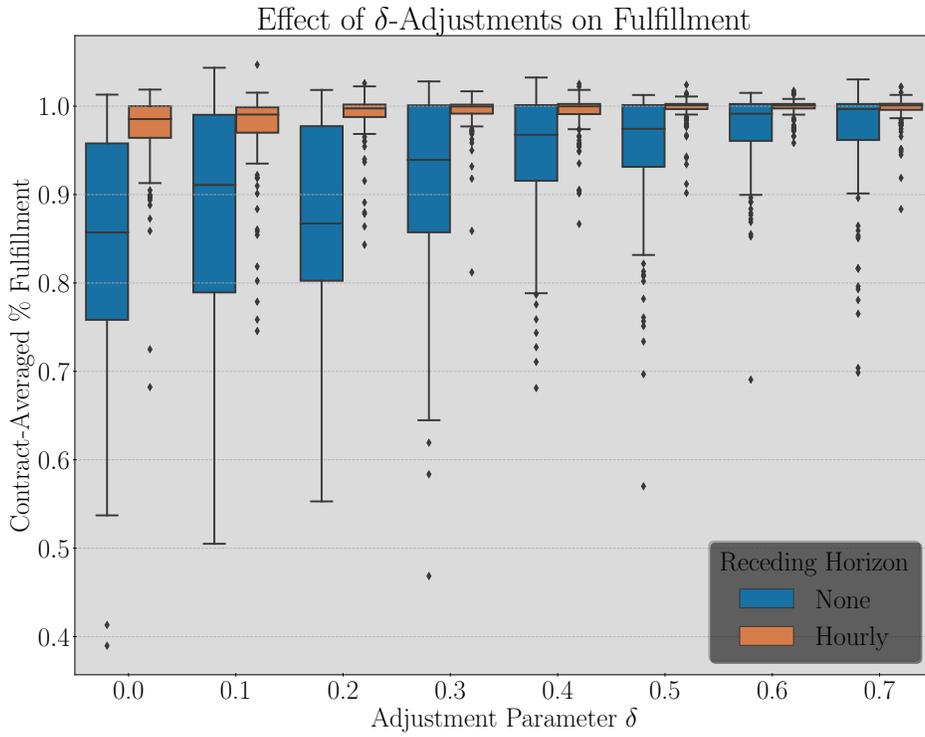
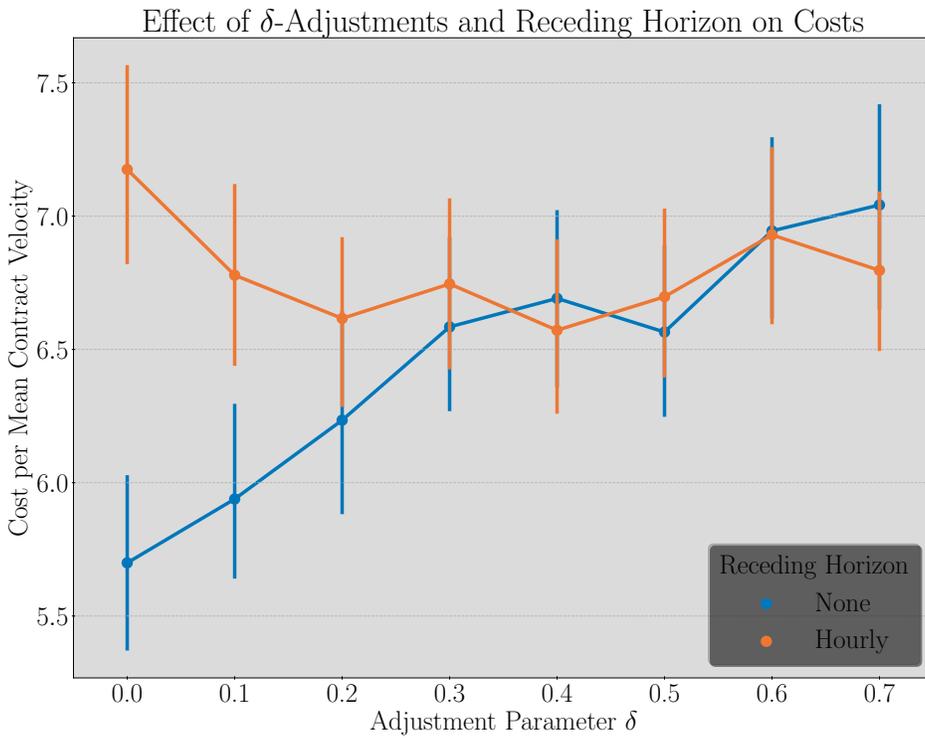Figure 2: Risk Adjustment Simulations (Fulfillment)

Figure 3: Risk Adjustment Simulations (Fulfillment)

simulation). An hourly receding horizon has a dramatic affect on the fulfillment – even without a risk adjustment, more than half of the contracts are fulfilled to at least 98%, and with a risk adjustment, almost all contracts are completely fulfilled.

Figure 3 provides a line plot similar to Figure 2 except the ordinate now quantifies the cost paid for items. Points are mean values across 500 simulations and error bars are 98% confidence intervals. In order to create a reasonable cost metric which is comparable across different contracts, we have normalized the total cost of attempting to fulfill a contract by that contract's "mean velocity" requirement. To understand what we refer to as the velocity requirement, suppose that we are obligated to obtain C items in T time. Then, we define the velocity of this requirement by $v = C/T$. We take the mean velocity of the whole contract to be the average of the velocities of each individual requirement in the contract (in this case, over the $N = 3$ requirements). To merely normalize by, e.g., the total number of items, would not result in a fair comparison since one contract may require obtaining the same number of items in less time (since they are generated randomly), which necessitates paying higher prices.

The main conclusions to be drawn from 3 are as follows. Firstly, without a receding horizon, there is naturally a tendency for costs to increase, since the supply acquisition is front-loaded. Secondly, as δ increases, the same upward pressure on costs is applicable as in the case without any receding horizon, but at the same time, for contracts which are difficult to fulfill, the receding horizon will drastically increase bids, which also increases the cost. However, for contracts which are easy to fulfill, or which get close to fulfillment early in the bidding period, the receding horizon smooths out the acquisition rate by reducing the bid, and therefore reducing costs. These are competing effects, but the use of the receding horizon is still able to reduce costs when δ is large; an ultimately small cost for achieving near complete fulfillment.

In summary, the receding horizon is able to achieve a much higher fulfillment proportion by increasing the bids (and incurring higher costs) for contracts which are difficult to fulfill (e.g., when supply rates are overestimated), while at the same time reducing the costs of fulfilling the remaining contracts by smoothing out bids that are too large.

4   Conclusions

This paper has studied the problem of minimum cost contract fulfillment in RTB auction markets. We have analyzed generalizations of the model proposed in [37] by considering both the first and second price auction mechanisms, and by allowing for general time-varying supply rate curves which can account for daily and weekly periodic cycles in market data. Moreover, we have uncovered the intrinsic convexity of the problem and provided a natural condition under which a strong duality result holds. Using duality, it has been shown that there is a close connection between the optimal bids $x(t)$ and certain dual variables which facilitates the complete representation of the optimal function $x(t)$ through the *finite* vector $\rho \in \mathbb{R}_+^N$ of dual variables. This close connection motivates the terminology "pseudo-bids" for the dual variables ρ. In the first price case, the optimal bids are obtained from ρ by mapping them through a time-varying function quantifying the marginal costs of acquiring additional items; in the second price case (but *not* the first), the optimal bids are given exactly by ρ, which implies the optimal bids in this case are piecewise constant functions of time.

Further consequences of duality have lead to the conclusion (backed up with experimental simulation evidence) that there is in fact very little difference between the optimal solution in the first price or second price cases. This result has the important practical implication that, regardless of the type of auction one is participating in, they may as well *assume* that they are participating in a second price auction, which has been seen to *always* lead to a convex problem. Indeed, in second price auctions, the optimal bids are *constant* (corresponding exactly to the dual pseudo-bids ρ) and therefore the original time-independent formulation of [37] largely covers the general time-dependent case after calculating suitable time averages.

Finally, Section 3 develops methods for implementation and reports the results of simulation studies. A discretization method for Problem (P) is studied in Section 3.1 and we obtain a convergence rate for this discretization. The discretized problem is used in an empirical Monte Carlo study in Section 3.2, which demonstrates the similarity between optimal bidding in first and second price auctions. Further simulations in Section 3.3 demonstrate the effectiveness of δ supply inflation and receding horizon bid recalculation for ensuring that contracts are actually fulfilled in practice. Additional information regarding the practical implementation of these methods is provided in the appendix, with the basic ideas being summarized in Algorithm 1.

References

[1]   A. Agrawal, R. Verschueren, S. Diamond, and S. Boyd. "A rewriting system for convex optimization problems". In: *Journal of Control and Decision* 5.1 (2018), pp. 42–60.

[2]   S. R. Balseiro, O. Besbes, and G. Y. Weintraub. "Repeated auctions with budgets in ad exchanges: Approximations and design". In: *Management Science* 61.4 (2015), pp. 864–884.

[3]   S. R. Balseiro and Y. Gur. "Learning in repeated auctions with budgets: Regret minimization and equilibrium". In: *Management Science* 65.9 (2019), pp. 3952–3968.

[4]   D. P. Bertsekas. "Extended monotropic programming and duality". In: *Journal of optimization theory and applications* 139.2 (2008), pp. 209–225.

[5]   S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. ISBN: 0521833787. URL: `https://web.stanford.edu/~boyd/cvxbook/`.

[6]   P. Bumpensanti and H. Wang. "A re-solving heuristic with uniformly bounded loss for network revenue management". In: *Management Science* 66.7 (2020), pp. 2993–3009.

[7]   H. Cai, K. Ren, W. Zhang, K. Malialis, J. Wang, Y. Yu, and D. Guo. "Real-Time Bidding by Reinforcement Learning in Display Advertising." In: *CoRR* abs/1701.02490 (2017). URL: `http://dblp.uni-trier.de/db/journals/corr/corr1701.html#CaiRZMWYG17`.

[8]   Y. Chen, P. Berkhin, B. Anderson, and N. R. Devanur. "Real-time bidding algorithms for performance-based display ad allocation." In: *KDD*. Ed. by C. Apté, J. Ghosh, and P. Smyth. ACM, 2011, pp. 1307–1315. ISBN: 978-1-4503-0813-7. URL: `http://dblp.uni-trier.de/db/conf/kdd/kdd2011.html#ChenBAD11`.

[9]   F. Clarke. *Functional analysis, calculus of variations and optimal control*. Vol. 264. Graduate Texts in Mathematics. Springer Science & Business Media, 2013. ISBN: 978-1-4471-4819-7. DOI: `10.1007/978-1-4471-4820-3`.

[10]   R. M. Corless, G. H. Gonnet, D. E. Hare, D. J. Jeffrey, and D. E. Knuth. "On the LambertW function". In: *Advances in Computational mathematics* 5.1 (1996), pp. 329–359.

[11]   Y. Cui, R. Zhang, W. Li, and J. Mao. "Bid landscape forecasting in online ad exchange marketplace." In: *KDD*. Ed. by C. Apté, J. Ghosh, and P. Smyth. ACM, 2011, pp. 265–273. ISBN: 978-1-4503-0813-7. URL: `http://dblp.uni-trier.de/db/conf/kdd/kdd2011.html#CuiZLM11`.

[12]   S. Diamond and S. Boyd. "CVXPY: A Python-embedded modeling language for convex optimization". In: *Journal of Machine Learning Research* 17.83 (2016), pp. 1–5.

[13]   J. Fernandez-Tapia, O. Guéant, and J.-M. Lasry. "Optimal real-time bidding strategies". In: *Applied Mathematics Research eXpress* 2017.1 (2017), pp. 142–183.

[14]   B. Gaujal, N. Navet, and C. Walsh. "Shortest-path algorithms for real-time scheduling of FIFO tasks with minimal energy use". In: *ACM Transactions on Embedded Computing Systems (TECS)* 4.4 (2005), pp. 907–933.

[15]   A. Ghosh, S. Mitra, S. Sarkhel, J. Xie, G. Wu, and V. Swaminathan. "Scalable Bid Landscape Forecasting in Real-time Bidding." In: *CoRR* abs/2001.06587 (2020). URL: `http://dblp.uni-trier.de/db/journals/corr/corr2001.html#abs-2001-06587`.

[16]   A. Ghosh, B. I. P. Rubinstein, S. Vassilvitskii, and M. Zinkevich. "Adaptive bidding for display advertising." In: *WWW*. Ed. by J. Quemada, G. León, Y. S. Maarek, and W. Nejdl. ACM, 2009, pp. 251–260. ISBN: 978-1-60558-487-4. URL: `http://dblp.uni-trier.de/db/conf/www/www2009.html#GhoshRVZ09`.

[17]   N. Grislain, N. Perrin, and A. Thabault. "Recurrent Neural Networks for Stochastic Control in Real-Time Bidding." In: *KDD*. Ed. by A. Teredesai, V. Kumar, Y. Li, R. Rosales, E. Terzi, and G. Karypis. ACM, 2019, pp. 2801–2809. ISBN: 978-1-4503-6201-6. URL: `http://dblp.uni-trier.de/db/conf/kdd/kdd2019.html#GrislainPT19`.

[18]   R. Gummadi, P. B. Key, and A. Proutière. "Optimal bidding strategies in dynamic auctions with budget constraints." In: *Allerton*. IEEE, 2011, p. 588. ISBN: 978-1-4577-1817-5. URL: `http://dblp.uni-trier.de/db/conf/allerton/allerton2011.html#GummadiKP11`.

[19]   IAB. *IAB internet advertising revenue report conducted by PricewaterhouseCoopers (PWC)*. 2018. URL: `https://www.iab.com/insights`.

[20]   K. Iyer, R. Johari, and M. Sundararajan. "Mean field equilibria of dynamic auctions with learning". In: *Management Science* 60.12 (2014), pp. 2949–2970.

[21]   N. Karlsson. "Adaptive control using Heisenberg bidding." In: *ACC*. IEEE, 2014, pp. 1304–1309. ISBN: 978-1-4799-3272-6. URL: `http://dblp.uni-trier.de/db/conf/amcc/acc2014.html#Karlsson14`.

[22]   N. Karlsson. "Control problems in online advertising and benefits of randomized bidding strategies". In: *European Journal of Control* 30 (2016), pp. 31–49.

[23] N. Karlsson. "Plant gain estimation in online advertising processes." In: *CDC*. IEEE, 2017, pp. 2182–2187.
      ISBN: 978-1-5090-2873-3. URL: `http://dblp.uni-trier.de/db/conf/cdc/cdc2017.html#Karlsson17`.

[24] V. Krishna. *Auction Theory*. Academic Press, 2002.

[25] W. H. Kwon and S. H. Han. *Receding horizon control: model predictive control for state models*. Springer
      Science & Business Media, 2006.

[26] L. J. Leblanc and L. Cooper. "The transportation-production problem". In: *Transportation Science* 8.4 (1974),
      pp. 344–354.

[27] H. Liao, L. Peng, Z. Liu, and X. Shen. "iPinYou Global RTB Bidding Algorithm Competition Dataset." In:
      ADKDD@KDD. Ed. by E. Saka, D. Shen, K. chih Lee, and Y. Li. ACM, 2014, 6:1–6:6. ISBN: 978-1-4503-
      2999-6. URL: `http://dblp.uni-trier.de/db/conf/kdd/adkdd2014.html#LiaoPLS14`.

[28] Y. Mansour, S. Muthukrishnan, and N. Nisan. "Doubleclick Ad Exchange Auction". In: *CoRR* abs/1204.0535
      (2012). URL: `http://dblp.uni-trier.de/db/journals/corr/corr1204.html#abs-1204-0535`.

[29] F. M. Menezes and P. K. Monteiro. *An introduction to auction theory*. OUP Oxford, 2005.

[30] A. Oberman. "The convex envelope is the solution of a nonlinear obstacle problem". In: *Proceedings of the
      American Mathematical Society* 135.6 (2007), pp. 1689–1694.

[31] A. B. Owen. *Monte Carlo theory, methods and examples*. 2013. URL: `http://statweb.stanford.edu/~owen/mc/`.

[32] C. Perlich, B. Dalessandro, R. Hook, O. Stitelman, T. Raeder, and F. J. Provost. "Bid optimizing and inventory
      scoring in targeted online advertising." In: *KDD*. Ed. by Q. Yang, D. Agarwal, and J. Pei. ACM, 2012, pp. 804–
      812. ISBN: 978-1-4503-1462-6. URL: `http://dblp.uni-trier.de/db/conf/kdd/kdd2012.html#PerlichDHSRP12`.

[33] A. Radovanovic and W. D. Heavlin. "Risk-aware revenue maximization in display advertising." In: *WWW*. Ed.
      by A. Mille, F. L. Gandon, J. Misselis, M. Rabinovich, and S. Staab. ACM, 2012, pp. 91–100. ISBN: 978-1-
      4503-1229-5. URL: `http://dblp.uni-trier.de/db/conf/www/www2012.html#RadovanovicH12`.

[34] M. H. Schneider. "Network flows and monotropic optimization". In: *Networks* 16.4 (1986), pp. 441–443. URL:
      `http://dblp.uni-trier.de/db/journals/networks/networks16.html#Schneider86`.

[35] J. F. Sharp, J. C. Snyder, and J. H. Greene. "A decomposition algorithm for solving the multifacility production-
      transportation problem with nonlinear production costs". In: *Econometrica: Journal of the Econometric Society*
      (1970), pp. 490–506.

[36] E. M. Stein and R. Shakarchi. *Real analysis: measure theory, integration, and Hilbert spaces*. Princeton Uni-
      versity Press, 2009.

[37] E. Tillberg, P. Marbach, and R. Mazumdar. "An Optimal Bidding Algorithm for Online Ad Auctions with
      Overlapping Targeting Criteria". In: *Proc. ACM Meas. Anal. Comput. Syst.* 4.2 (June 2020). DOI: `10.1145/3366707`.

[38] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W.
      Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson,
      E. Jones, R. Kern, E. Larson, C. Carey, İ. Polat, Y. Feng, E. W. Moore, J. Vand erPlas, D. Laxalde, J. Perktold,
      R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van
      Mulbregt, and S. 1. 0. Contributors. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python".
      In: *Nature Methods* 17 (2020), pp. 261–272.

[39] J. Wang, W. Zhang, and S. Yuan. "Display advertising with real-time bidding (RTB) and behavioural targeting".
      In: *Foundations and Trends® in Information Retrieval* 11.4-5 (2017), pp. 297–435.

[40] Y. Wang and H. Wang. "Constant Regret Re-solving Heuristics for Price-based Revenue Management". In:
      arXiv preprint arXiv:2009.02861 (2020).

[41] Y. Wang, K. Ren, W. Zhang, J. Wang, and Y. Yu. "Functional Bid Landscape Forecasting for Display Ad-
      vertising." In: *ECML/PKDD (1)*. Ed. by P. Frasconi, N. Landwehr, G. Manco, and J. Vreeken. Vol. 9851.
      Lecture Notes in Computer Science. Springer, 2016, pp. 115–131. ISBN: 978-3-319-46127-4. URL: `http://dblp.uni-trier.de/db/conf/pkdd/pkdd2016-1.html#WangRZWY16`.

[42] L. Wasserman. *All of nonparametric statistics*. Springer Science & Business Media, 2006. DOI: `10.1007/0-387-30623-4`.

[43] D. Wu, X. Chen, X. Yang, H. Wang, Q. Tan, X. Zhang, J. Xu, and K. Gai. "Budget constrained bidding by
      model-free reinforcement learning in display advertising". In: *Proceedings of the 27th ACM International Con-
      ference on Information and Knowledge Management*. 2018, pp. 1443–1451.

[44] W. C.-H. Wu, M.-Y. Yeh, and M.-S. Chen. "Predicting Winning Price in Real Time Bidding with Censored
      Data." In: *KDD*. Ed. by L. Cao, C. Zhang, T. Joachims, G. I. Webb, D. D. Margineantu, and G. Williams. ACM,
      2015, pp. 1305–1314. ISBN: 978-1-4503-3664-2. URL: http://dblp.uni-trier.de/db/conf/kdd/
      kdd2015.html#WuYC15.

[45] W. C.-H. Wu, M.-Y. Yeh, and M.-S. Chen. "Deep Censored Learning of the Winning Price in the Real Time
      Bidding." In: *KDD*. Ed. by Y. Guo and F. Farooq. ACM, 2018, pp. 2526–2535. URL: http://dblp.uni-
      trier.de/db/conf/kdd/kdd2018.html#WuYC18.

[46] S. Yuan, J. Wang, and X. Zhao. "Real-time bidding for online advertising: measurement and analysis". In:
      Proceedings of the Seventh International Workshop on Data Mining for Online Advertising. 2013, pp. 1–8.

[47] W. Zhang, Y. Rong, J. Wang, T. Zhu, and X. Wang. "Feedback Control of Real-Time Display Advertising."
      In: *WSDM*. Ed. by P. N. Bennett, V. Josifovski, J. Neville, and F. Radlinski. ACM, 2016, pp. 407–416. ISBN:
      978-1-4503-3716-8. URL: http://dblp.uni-trier.de/db/conf/wsdm/wsdm2016.html#ZhangRWZW16.

[48] W. Zhang and J. Wang. "Statistical Arbitrage Mining for Display Advertising." In: *KDD*. Ed. by L. Cao, C.
      Zhang, T. Joachims, G. I. Webb, D. D. Margineantu, and G. Williams. ACM, 2015, pp. 1465–1474. ISBN:
      978-1-4503-3664-2. URL: http://dblp.uni-trier.de/db/conf/kdd/kdd2015.html#ZhangW15a.

[49] W. Zhang, S. Yuan, and J. Wang. "Optimal real-time bidding for display advertising." In: *KDD*. Ed. by S. A.
      Macskassy, C. Perlich, J. Leskovec, W. Wang, and R. Ghani. ACM, 2014, pp. 1077–1086. ISBN: 978-1-4503-
      2956-9. URL: http://dblp.uni-trier.de/db/conf/kdd/kdd2014.html#ZhangYW14.

[50] W. Zhang, S. Yuan, J. Wang, and X. Shen. "Real-time bidding benchmarking with ipinyou dataset". In: *arXiv
      preprint arXiv:1407.7073* (2014).

[51] W. Zhang, T. Zhou, J. Wang, and J. Xu. "Bid-aware Gradient Descent for Unbiased Learning with Censored
      Data in Display Advertising." In: *KDD*. Ed. by B. Krishnapuram, M. Shah, A. J. Smola, C. C. Aggarwal,
      D. Shen, and R. Rastogi. ACM, 2016, pp. 665–674. ISBN: 978-1-4503-4232-2. URL: http://dblp.uni-
      trier.de/db/conf/kdd/kdd2016.html#ZhangZWX16.

[52] W.-Y. Zhu, W.-Y. Shih, Y.-H. Lee, W.-C. Peng, and J.-L. Huang. "A gamma-based regression for winning
      price estimation in real-time bidding advertising." In: *BigData*. Ed. by J.-Y. Nie, Z. Obradovic, T. Suzumura,
      R. Ghosh, R. Nambiar, C. Wang, H. Zang, R. Baeza-Yates, X. Hu, J. Kepner, A. Cuzzocrea, J. Tang, and M.
      Toyoda. IEEE Computer Society, 2017, pp. 1610–1619. ISBN: 978-1-5386-2715-0. URL: http://dblp.uni-
      trier.de/db/conf/bigdataconf/bigdataconf2017.html#ZhuSLPH17.

A Supply Rate Curve Estimation

In practice, the functions $W(x)$ will be estimated from available historical data. Therefore, depending on the method used to carry out this estimation, the resulting $\Lambda(x)$ is not necessarily guaranteed to be convex in the first price case (recall Proposition 2.2 requires 2-concavity of W). If the estimate of $W(x)$ is simply carried out by fitting a particular parameterized distribution (e.g., a normal approximation) to a dataset, then there is unlikely to be any issue since most distributions commonly employed for this purpose do in fact have log-concave cumulative distribution functions. Moreover, since $W(x)$ arises through an auction process, extreme value theory and the Fisher-Tippett-Gnedenko Theorem may motivate the belief that $W(x)$ should tend to be close to a Weibull distribution, which is log-concave at least for some parameter ranges.

However, our empirical data (see also [27, 50]) suggests that such simple models are *not* good estimates for supply rate curves, and that the curves have a tendency towards some multi-modality. For this reason, the methods of Section 3.2 make use of Kernel density estimates (KDE). A Gaussian KDE model for $W(x)$ is natural because it ensures smoothness, that $W > 0$, and the bandwidth can be chosen to correspond to the level of bid noise (or the level of bid noise can be chosen from the optimal KDE bandwidth).

Unfortunately, KDE estimates of $W(x)$ from data need not be (and often aren't) log-concave or α-concave. In order to deal with this problem, we consider calculating a convex and piecewise affine majorant of the function $\Lambda(s)$, which will transform Problem (P_K) into a linear program. A similar minorizing envelope can also be calculated through the methods of [30], and a piecewise minorant thereof computed by outer linear approximations. There is no obvious reason to prefer one approximation over the other and our experience has not demonstrated a clear benefit either way.

Alternatively, taking the log-concave envelope of $W(x)$ through similar methods is guaranteed to result in convex acquisition functions. The considerations of the previous paragraph suggest that we should at least expect the KDE estimates of W to be "almost" log-concave, and indeed, this is what we have observed in our own experiments; the convex envelopes are only slight perturbations of the original supply rate curve estimate.

A.1 Piecewise Affine Approximation

Let us denote by $\widetilde{\Lambda}(x)$ the acquisition cost function attained from an estimated supply rate curve and by $\Lambda^{\mathsf{U}}(x)$ the minimal convex majorant $\widetilde{\Lambda}(x)$, i.e.,

$$\Lambda^{\mathsf{U}}(x) = \inf\{\lambda(x) \mid \lambda(x) \geq \widetilde{\Lambda}(x), \lambda \text{ is convex and monotone increasing}\}. \tag{15}$$

We emphasize that λ must be monotone increasing, but this would also follow as a consequence of the monotonicity of $\widetilde{\Lambda}$. The maximal minorant can be defined similarly. Moreover, α-concave envelopes can be calculated by requiring that $\ell_\alpha \circ \lambda$ is convex.

A piecewise affine approximation of Λ^{U} can be found by discretizing a compact interval $[a, b] \subset \mathbb{R}$ into $n + 1$ points $x_0, x_1, \ldots, x_n$ and solving the following convex quadratic program where convexity and monotonicity are enforced via finite differences

$$
\begin{aligned}
\underset{\lambda}{\text{minimize}} \quad & \frac{1}{n+1} \sum_{i=0}^{n} \left(\lambda_i - \widetilde{\Lambda}(x_i)\right)^2 \\
\text{subject to} \quad & \lambda_i \geq \widetilde{\Lambda}(x_i) \\
& \lambda_i - \lambda_{i-1} \geq 0 \; \forall i \in [n] \\
& \lambda_{i+1} - 2\lambda_i + \lambda_{i-1} \geq 0 \; \forall i \in [n-1].
\end{aligned}
\tag{E^{U}}
$$

An accurate approximation of the convex majorant is recovered via linearly interpolating λ_i. In fact, λ_i will result in a *strictly* monotone function (and therefore a continuous inverse) whenever $\widetilde{\Lambda}$ is strictly monotone.

It is important that the developments in Section 2.4 do not make any assumptions regarding the differentiability of $W(x)$ or $\Lambda(s)$, since if these curves are corrected through solution of Problem (E^{U}) to ensure convexity, it is by definition not differentiable. Moreover, simply interpolating the resulting λ_i (e.g. with a cubic spline) may not be acceptable as, to our knowledge, it is not possible for this process to maintain simultaneously the monotonicity, convexity, and minorization of $\widetilde{W}$.

Remark A.1 (Sparse Approximations). It is desirable to use a fine discretization in Problem (E^U), otherwise the resulting function may fail to majorize $\widetilde{\Lambda}$ in regions of high curvature. However, each piecewise segment translates to an additional constraint when Λ^U is substituted into Problem (P_K) (see Section A.3), which may become overly burdensome for large problems. Therefore, it may be desirable to use a coarser approximation obtained by linearly interpolating samples of $\Lambda^U(s)$. Since $\Lambda^U(s)$ is convex, this process is guaranteed to produce another piecewise affine convex function which further majorizes $\Lambda^U(s)$.

A.2   An Example

We consider an illustrative example of calculating log-concave envelopes of the supply rate curve in a simple market model. We suppose that each participant is characterized by a bid and rate pair (b_i, r_i) indicating that they will bid b_i with probability r_i on any arriving item. We sample $\{(b_i, r_i)\}_{i=1}^{30}$ randomly as $b \sim \exp(0.5)$ and $r_i \sim \beta(11.1, 10)$ which represents a market with 30 participants whose average bid is 0.5 and have an average probability of $(11.1 - 10)/11.1 \approx 0.1$ of bidding.

The *bid landscape* in this situation is given by

$$W(x) = \prod_{i:b_i > x} (1 - r_i),$$

indicating the probability of winning an item if the bid x is placed.

We let $\widetilde{W}(x)$ be a KDE smoothed (with $\sigma^2 = 1/4$) version of W which corresponds either to the true supply rate curve under randomized bidding, or a reasonable estimate (from historical data) thereof. We denote supply rate curve estimates $\widetilde{W}^L(x)$ and $\widetilde{W}_{\mathcal{N}}$ which are derived from $W(x)$ by solving Problem (15) for the function $-\log \circ W$ and moment matching a Gaussian c.d.f., respectively. Note that this procedure produces *minorants* of the supply rate curve (and therefore majorants of the acquisition cost curve), since using a *convex* majorant procedure results in minorants of *concave* functions. Figure 4 plots examples of these functions and their associated cost and acquisition counterparts.

The examples of Figure 4 are chosen to deliberately exaggerate the differences between the supply curve estimates and the envelopes. For the simulation examples of Section 3.2 the two curves are often indistinguishable.

A.3   Linear Approximations of Primal Problem

From Section 3.1 we have the time discretized primal problem (P_K):

$$
\begin{aligned}
\underset{s,r}{\text{minimize}} \quad & \sum_{j=1}^{M} \sum_{k:\widetilde{T}_k < T^j} \overline{\Lambda}_{jk}(s_j[k]) \\
\text{subject to} \quad & \sum_{j \in \mathcal{A}_i} \sum_{k:\widetilde{T}_k < T^j} \Delta_k r_{ij}[k] \geq C_i \\
& \sum_{i \in \mathcal{B}_j \cap \mathcal{T}_{T_k}} r_{ij}[k] = s_j[k] \\
& r_{ij}[k] \geq 0.
\end{aligned}
\qquad (P_K)
$$

Suppose that the functions $\Lambda_j(s, \widetilde{T}_k)$ have piecewise affine approximations $\widehat{\Lambda}_k(s) = \max_{h \in H_{jk}} \left(m_{hj}[k]s + b_{hj}[k]\right)$. Then, by introducing additional variables α_{jk} we can reformulate P_K in epigraph form as a linear program
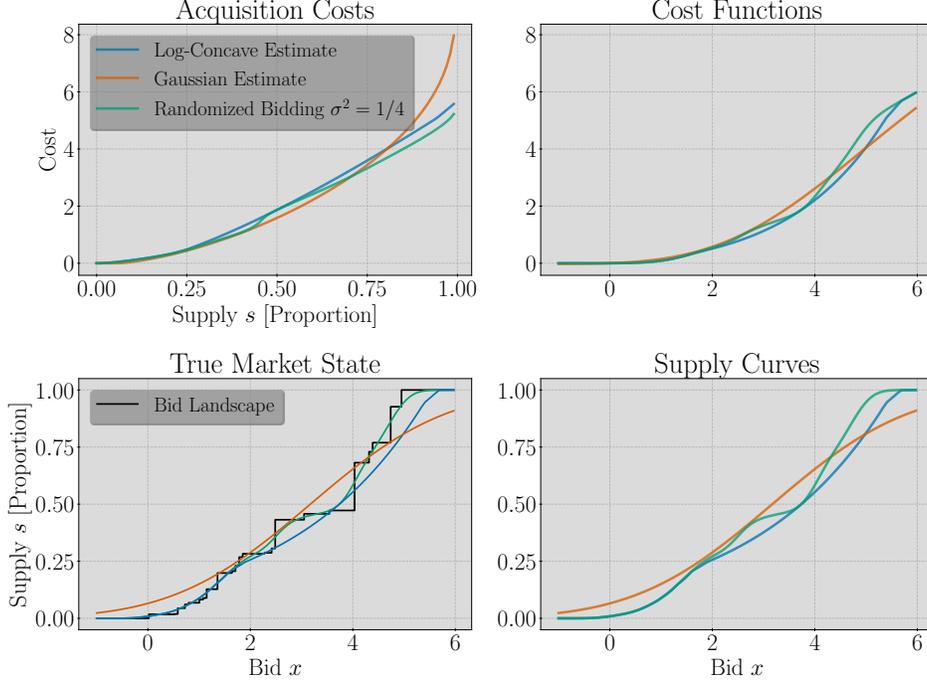
Figure 4: Example Supply Rate Functions

Comparison of different methods of estimating supply rate curves. Lower left: Comparison of KDE smoothing, the maximal log-concave minorant thereof, and a Gaussian c.d.f. (fit by moment matching) overlaid upon a true market state. Lower right: The three corresponding supply rate curves. Upper right: Corresponding cost curves. Upper left: Corresponding acquisition cost functions where we see that KDE smoothing does not lead to convexity, and that a Gaussian estimate is not a consistent minorant or majorant.

$$
\begin{aligned}
\underset{s,r}{\text{minimize}} \quad & \frac{1}{2} \sum_{j=1}^{M} \sum_{k:\widetilde{T}_k \leq T^j} \Delta_k (\alpha_{jk} + \alpha_{j,k-1}) \\
\text{subject to} \quad & \sum_{j \in \mathcal{A}_i} \sum_{k:\widetilde{T}_k \leq T^j} \Delta_k r_{ij}[k] \geq C_i \\
& m_{hj}[k] s_j[k] + b_{hj}[k] \leq \alpha_{jk} \\
& \sum_{i \in \mathcal{B}_j \cap \mathcal{T}_{T_k}} r_{ij}[k] = s_j[k] \\
& r_{ij}[k] \geq 0,
\end{aligned}
\tag{16}
$$

which is the formulation we have employed for our simulations in Section 3.2.

B Simulation (Additional Details)

In this section we provide additional details on the methods used to produce the results of Section 3.2.

B.1 Estimating Supply Rate Curves

The iPinYou dataset consists of impression data derived from a real DSP and includes information about bidding prices, market prices, and user characteristics. We focus on the *season two* data (a week long period 2013-06-06 to 2013-06-12). In all cases, our supply rate curve estimates are 24h-periodic in time and therefore account for the natural daily (but not weekly) cycles in prices and arrival rates. Since this paper does not focus on the estimation of supply rate curves, we apply a simple estimation procedure using the entire dataset as input. Though this has the
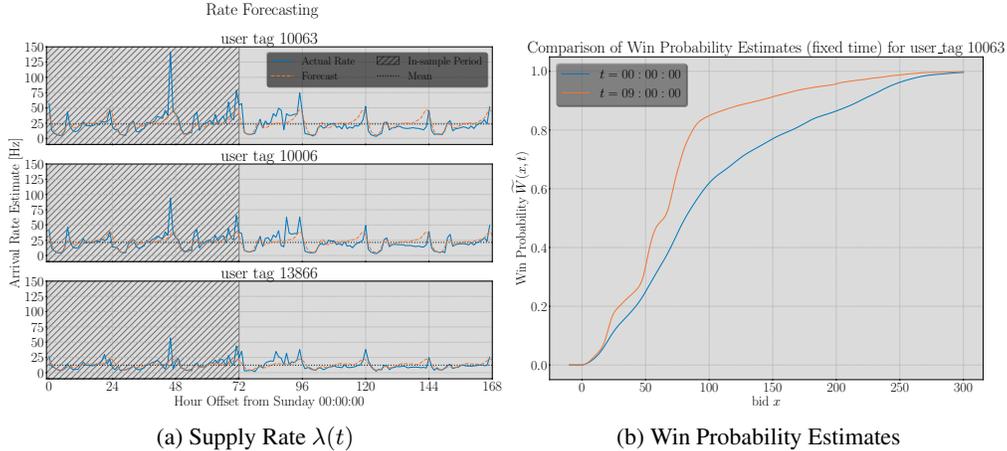
(a) Supply Rate $\lambda(t)$ (b) Win Probability Estimates

Figure 5: Estimated Supply Rate Curves and Costs
Cost and supply rate curves estimated from iPinYou data from the first 3 days of season 2. (a) Item arrival rates and
the corresponding forecasts. The hatched region indicates an in-sample period with the remainder being
out-of-sample. Our simulations run on similar 3 day periods with a 12 hour sliding window for a total of 9 periods of
72 hours each. (b) Estimates of the win probability function $\widetilde{W}(x,t)$ for $t = 00:00:00$ (blue) and $t = 09:00:00$
(red). We compare Gaussian KDE (solid line) with a parametric Exponential CDF (dashed line).

effect of leaking some information from the future, the estimation of supply rate curves is not subject to optimization, limiting the impact of this leakage. Moreover, the dataset is averaged into a single 24h periodic function and extended through periodicity. It is reasonable to believe that the previous week's (out of sample) data would provide similar results. Figure 5 provides an illustration of estimated supply rate curves where only 72h of data is used to forecast the remaining 96h for purposes of illustrating a case where historical data is used in this manner to forecast future trends.

Provided with the iPinYou dataset is a `user_tag` which, according to [50] is "[a segment] in iPinYou's proprietary audience database". We therefore use the `user_tag` property as the "item types", focusing on the five most common tags: `10063`, `10006`, `13866`, `10024`, and `10083`.

B.2   Estimating the Supply Rate

In order to estimate the supply rate $\lambda_j(t)$ for each user tag, we have taken the inverse of the average of the time differences between arrival instants in each hour of the day, after removing outliers. Calculating the number of arrivals over an hour long period is not adequate as there appear to be large consistent gaps in arrival times: we suspect that the dataset was subsampled prior to being released.

This calculation results in estimates $\lambda_j[0], ..., \lambda_j[23]$, with time denoted in hours. The continuous estimate was subsequently formed by smoothly interpolating between these points with a $24-$periodic boundary, resulting in a function $\tilde{\lambda}_j(t)$ defined on $[0, 24]$. A forecast for the average supply rate at time t is obtained via $\lambda_j(t) = \tilde{\lambda}(t \bmod 24)$.

An illustrative example for tags `10063`, `10006`, and `13866` is provided in Figure 5a.

B.3   Estimating Win Probabilities

Similarly to the supply rate estimates, we estimate an average win probability function for each $t \in 0, \dots, 23$ and then smoothly interpolate along t to estimate a $24-$periodic function $\widetilde{W}_j(x,t)$ indicating the probability of winning an impression of type (user tag) j arriving at time t given a bid x.

The estimate of $x \mapsto \widetilde{W}_j(x,t)$ is obtained by smoothing the histogram with a Gaussian kernel (bandwidth chosen simply by the Normal Reference Rule [42, Chap. 6.3]) for each `market_price` data point falling into the hour long window. The results of this procedure, as well as a comparison to a parametric estimate with an Exponential density are given in Figure 5b.

The `market_price` attribute in the dataset corresponds to the price actually paid in the second price auction. We have not accounted for the affects of censoring; since the DSP collected the dataset with large bids intended to win most impressions that were bid on, this isn't a significant factor.

23

B.4   Cost and Supply Rate Curves

The supply rate curve $W_j(x,t)$ is simply the product of the supply rate $\lambda_j(t)$ and the win probability $\widetilde{W}_j(x,t)$. The cost curves f^{1st} and f^{2nd} is derived from the supply rate curve using numerical integration and interpolation.

B.5   Simulating the Bidding Process

The simulations of Section 3.2 are obtained by storing the hour-by-hour inter-arrival and price data for each item type $j \in [M]$ and sampling uniformly from these datasets. At simulation time $t \in \mathbb{R}_+$ we sample an inter-arrival time Δt and price P from the data for hour $\lfloor t \rfloor + 1$ with probability $t - \lfloor t \rfloor$ and otherwise from the data for hour $\lfloor t \rfloor$. A bid is solicited from a bidder (an implementation of (P_K)) and if the bid exceeds P the bidder allocates that item to the fulfillment of a contract. The simulation time is them updated to $t + \Delta t$ and the process continues.

Algorithm 2: Bidding Simulation

---

input   : A Bidder derived from Section 3.1.
output : Recording of Bidder's item allocations to process into normalized acquisition curves.

1 // Initialize:
2 $Q \leftarrow$ Priority-Queue($[\,]$) // Sort by time
3 $t \leftarrow 0$ // The ``current'' time
4 **for** $j \in [M]$ **do**
5   // Sample an inter-arrival time and a price
6      $(\Delta t, P) \leftarrow$ Sample-Dataset(t,j)
7   Q.push($(t + \Delta t, P, j)$)

8 // Simulate bidding process:
9 **while** $t < T_{end}$ **do**
10     $t, P, j \leftarrow Q$.pop()
11     $b \leftarrow$ Bidder.solicit_bid(t,j) // Ask for a bid on type j at time t
12   **if** $b \geq P$ **then**
13   Bidder.award_item(t,j) // Allocate items for winning bids
14     $(\Delta t, P) \leftarrow$ Sample-Dataset(t,j) // Append next (t, P) pair for j to Q
15   Q.push($(t + \Delta t, P, j)$)

16 **Function** Sample-Dataset(t,j):
17     $p \leftarrow t - \lfloor t \rfloor$
18     $U \sim \mathcal{U}(0,1)$ // Interpolate between hours
19   **if** $p \leq U$ **then**
20        $h \leftarrow \lfloor t \rfloor$
21   **else**
22        $h \leftarrow \lfloor t \rfloor + 1$
23     $\Delta t \leftarrow$ Sample-Inter-arrivals(hour=h,type=j)
24     $P \leftarrow$ Sample-Prices(hour=h,type=j)
25   **return** $(\Delta t, P)$

C   Proofs

C.1   Proofs of Main Results

Proposition 2.2 (Convex Acquisition Costs – First Price Case). *Suppose that a supply rate curve $W(x)$ is strictly 2-concave. Then in a first price auction, the extended acquisition function*

$$\Lambda^{1st}(s) \triangleq \begin{cases} \infty; & s > \bar{s} \\ 0; & s < W(0) \\ sW^{-1}(s); & \text{otherwise} \end{cases}, \tag{6}$$

is a proper[3], lower semi-continuous, non-decreasing, and convex function on $\mathbb{R}$, which is strictly convex on $[W(0), \bar{s}]$. Moreover, the convex conjugate

$$\Lambda_{1st}^*(p) \stackrel{\Delta}{=} \sup_{s \in \mathbb{R}} [sp - \Lambda(s)],$$

is itself a proper, lower semi-continuous, non-decreasing, and convex function on $\mathbb{R}$, which satisfies $\Lambda_{1st}^*(p) = \infty$ for $p < 0$, and if $\bar{s} < \infty$, then $\Lambda_{1st}^*(p) = p\bar{x} - \Lambda_{1st}(\bar{s})$ on $p > \bar{x}$.

Proof of Proposition 2.2. Since $f(x) = xW(x)\mathbf{1}_{\mathbb{R}_+}(x)$ we have, for $s \in \text{ran } W$, $\Lambda(s) = sW^{-1}(s)$. On $s \leq W(0)$ we have $\Lambda(s) = 0$, and on $s \in \{s | s > \bar{s}\}$ we have $\Lambda(s) = \infty$. Moreover, by Assumption 2.1, W attains its supremum (when finite) and so Λ is continuous at $\bar{s}$. Convexity will therefore follow if Λ is convex on int ran W. To this end, we use the 2-concavity of W to see that $1 - 1/W(x)$ is concave on its domain and therefore that the inverse, $W^{-1}\big(1/(1-x)\big)$ is convex for the portion of $x \in (-\infty, 1)$ which remains in the range of W. It is fairly well known that for a convex function f, the function

$$(cx + d)f\Big(\frac{ax + b}{cx + d}\Big)$$

is convex on $cx + d > 0$ (see e.g. [5, Ex. 3.20]). Therefore, by setting $a = c = 1$, $b = -1$ and $d = 0$ we obtain convexity of

$$(cs + d)W^{-1}\Big(\frac{cs + d}{(c - a)s + (d - b)}\Big) = sW^{-1}(s)$$

which is the function $\Lambda(s)$.

In consideration of the conjugate, it is evident that since $\Lambda(s) \geq 0$ we have $\Lambda^*(p) = \infty$ for $p < 0$ (take $s \to -\infty$). For $p > \bar{x}$, if $\bar{s} < \infty$, then since $s < \bar{s} \implies W^{-1}(s) < \bar{x}$ and $\Lambda(\bar{s}+) = \infty$ we see that $sp - \Lambda(s)$ is maximized at $s = \bar{s}$. That $\Lambda^*(p)$ is non-decreasing follows since for $p \leq q$ we have $sp - \Lambda(s) \leq sq - \Lambda(s)$ on $s \geq 0$. Finally, that $\Lambda^*(p)$ is proper, lower semi-continuous, and convex is a statement of the Fenchel-Moreau Theorem (see e.g., [9, sec 4.2]).

$\square$

Proposition 2.3 (Continuous Primal Problem). *In a first or second price auction, suppose that for each $j \in [M]$ and $t \in [0, T^j]$, the acquisition cost curve $\Lambda_j(x, t)$ is convex in x. Then, Problem (TP) can be reformulated as*

$$
\begin{aligned}
\underset{s,r}{\text{minimize}} \quad & \sum_{j=1}^{M} \int_0^{T^j} \Lambda_j(s_j(t), t)\mathrm{d}t \\
\text{subject to} \quad & \sum_{j \in \mathcal{A}_i} \int_0^{T_i} r_{ij}(t)\mathrm{d}t \geq C_i \\
& \sum_{i \in \mathcal{B}_j} r_{ij}(t) = s_j(t) \\
& r_{ij}(t) \geq 0.
\end{aligned}
\qquad (P)
$$

If a solution exists, then a solution to the original problem (TP) is obtained via $x_{ij}(t) = W_j^{-1}(s_j(t), t)$ for each $i \in \mathcal{B}_j$ and $\gamma_{ij}(t) = r_{ij}(t)/s_j(t)$. Moreover, Problem (P) is convex.

Proof of Proposition 2.3. First, we establish that whenever a solution exists, there is also a solution with the property that $x_{ij}(t) = x_j(t)$ and where $\sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) \in \{0, 1\}$. To this end, suppose (x, γ) is a solution of Problem (TP), and

---

[3]Recall that a function $f : \mathbb{R} \to (-\infty, \infty]$ is *proper* if it is not everywhere equal to $+\infty$.

with total cost J. Let $(\tilde{x}, \tilde{\gamma})$ be another bid and allocation pair with total cost $\tilde{J}$ defined by

$$\tilde{x}_j(t) \triangleq W_j^{-1}\Big(\sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) W_j(x_{ij}(t), t), t\Big),$$

$$\tilde{\gamma}_{ij}(t) \triangleq \frac{\gamma_{ij}(t) W_j(x_{ij}(t), t)}{\sum_{u \in \mathcal{B}_j} \gamma_{uj}(t) W_j(x_{uj}(t), t)},$$

where $0/0 \triangleq 0$ in the definition of $\tilde{\gamma}$. We proceed to show that $(\tilde{x}, \tilde{\gamma})$ is also a solution and we note that the definition of $\tilde{\gamma}$ satisfies $\forall t < T^j \ \sum_{i \in \mathcal{B}_j} \tilde{\gamma}_{ij}(t) \in \{0, 1\}$.

It is clear that the pair $\tilde{x}, \tilde{\gamma}(t)$ is feasible by construction. Indeed, $\tilde{\gamma}_{ij}(t) \geq 0$ and $\sum_{i \in \mathcal{B}_j} \tilde{\gamma}_{ij}(t) \leq 1$ by definition. Moreover, we have

$$\sum_{j \in \mathcal{A}_i} \tilde{\gamma}_{ij}(t) W_j(\tilde{x}_j(t), t) = \sum_{j \in \mathcal{A}_i} \Big[\frac{\gamma_{ij}(t) W_j(x_{ij}(t), t) \sum_{v \in \mathcal{B}_j} \gamma_{vj}(t) W_v(x_{vj}(t), t)}{\sum_{u \in \mathcal{B}_j} \gamma_{uj}(t) W_j(x_{uj}(t), t)} \Big]$$

$$= \sum_{j \in \mathcal{A}_i} \gamma_{ij}(t) W_j(x_{ij}(t), t),$$

which when integrated from 0 to T_i meets or exceeds C_i since (x, γ) is assumed to be a solution.

The cost of $(\tilde{x}, \tilde{\gamma})$, instantaneously at time t, then satisfies $\tilde{J} = J$ since J is the minimal cost and

$$\tilde{J} \triangleq \int_0^T \Big[\sum_{i=1}^N \sum_{j \in \mathcal{A}_i} \tilde{\gamma}_{ij}(t) f_j(\tilde{x}_j(t)) \Big] \mathrm{d}t$$

$$\stackrel{(a)}{=} \int_0^T \Big[\sum_{i=1}^N \sum_{j \in \mathcal{A}_i} \tilde{\gamma}_{ij}(t) \Lambda_j \Big(\sum_{u \in \mathcal{B}_j} \gamma_{uj}(t) W_j(x_{uj}(t), t), t \Big) \Big] \mathrm{d}t$$

$$\stackrel{(b)}{\leq} \int_0^T \Big[\sum_{i=1}^N \sum_{j \in \mathcal{A}_i} \tilde{\gamma}_{ij}(t) \sum_{u \in \mathcal{B}_j} \gamma_{uj}(t) \Lambda_j(W_j(x_{uj}(t), t), t) \Big] \mathrm{d}t$$

$$\stackrel{(c)}{=} \int_0^T \Big[\sum_{j=1}^M \sum_{u \in \mathcal{B}_j} \gamma_{uj}(t) f_j(x_{uj}(t), t) \sum_{i \in \mathcal{B}_j} \tilde{\gamma}_{ij}(t) \Big] \mathrm{d}t = J$$

where (a) is just the definition of $\Lambda_j = f_j \circ W_j^{-1}$ (see Section 2.3), (b) follows by the convexity of Λ_j and that $\Lambda_j(0) = 0$ (since γ_{ij} need not necessarily sum to 1), and (c) follows again by $\Lambda_j = f_j \circ W_j^{-1}$ and then by swapping the order of summation using $i \in \mathcal{B}_j \iff j \in \mathcal{A}_i$.

We can now recall the original problem (TP), and apply the preceding results to eliminate the dependence of the bid on i, since if a solution exists it can be assumed to have the property $x_{ij}(t) = x_j(t)$:

$$\underset{x, \gamma}{\text{minimize}} \quad \sum_{i=1}^N \int_0^{T_i} \Big[\sum_{j \in \mathcal{A}_i} \gamma_{ij}(t) f_j(x_j(t), t) \Big] \mathrm{d}t$$

$$\text{subject to} \quad \sum_{j \in \mathcal{A}_i} \int_0^{T_i} \gamma_{ij}(t) W_j(x_j(t), t) \mathrm{d}t \geq C_i$$

$$\sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) \leq 1, \gamma_{ij}(t) \geq 0,$$

Due to the bids' independence of i, we can rearrange the objective by swapping the order of summation:

$$\sum_{i=1}^{N} \int_0^{T_i} \Big[\sum_{j \in \mathcal{A}_i} \gamma_{ij}(t) f_j(x_j(t), t) \Big] \mathrm{d}t = \sum_{j=1}^{M} \int_0^{T^j} \Big[f_j(x_j(t), t) \sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) \Big] \mathrm{d}t, \tag{17}$$

which, after making the substitution $s_j(t) = W_j(x_j(t), t)$, results in

$$\begin{aligned}
\underset{s,\gamma}{\text{minimize}} \quad & \sum_{j=1}^{M} \int_0^{T^j} \Big[\Lambda_j(s_j(t), t) \sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) \Big] \mathrm{d}t \\
\text{subject to} \quad & \int_0^{T_i} \sum_{j \in \mathcal{A}_i} \gamma_{ij}(t) s_j(t) \geq C_i \\
& \sum_{i \in \mathcal{B}_j \cap \mathcal{T}_t} \gamma_{ij}(t) \leq 1 \\
& \gamma_{ij}(t) \geq 0, s_j(t) \geq 0.
\end{aligned} \tag{18}$$

Now, since any optimal solution can be assumed to be such that $\sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) \in \{0, 1\}$ and since $\Lambda_j(0, t) = 0$, there is no loss of generality in adding the constraint $\sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) = 1$ and eliminating this term from the objective. Finally, making the substitution $r_{ij}(t) = \gamma_{ij}(t) s_j(t)$ results in the formulation of Problem (P). □

Proposition 2.4 (Duality). *A Dual of (P) can be formulated as*

$$\begin{aligned}
\underset{\rho,\mu}{\text{maximize}} \quad & -\sum_{j=1}^{M} \sum_{k:T_k \leq T^j} \int_{T_{k-1}}^{T_k} \Lambda_j^*(\mu_{jk}, t) \mathrm{d}t + \sum_{i=1}^{N} \rho_i C_i \\
\text{subject to} \quad & \mu_{jk} \geq \rho_i \; \forall i \in \mathcal{B}_j \cap \mathcal{T}_{T_k} \\
& \rho_i \geq 0,
\end{aligned} \tag{D}$$

which is a finite convex problem. Problem (D) is dual to Problem (P) in the sense that if $D^\star$ and $P^\star$ are their respective values (possibly ∞ or $-\infty$), then $D^\star \leq P^\star$.

Finally, under Assumption 2.2 there exists a solution $(s, r) \in L_2(\mathbb{R})^M \times L_2(\mathbb{R})^{N \times M}$ to Problem (P) and a solution $(\rho, \mu) \in \mathbb{R}^N \times \mathbb{R}^M$ to Problem (D) and $-\infty < D^\star = P^\star < \infty$.

Proof of Proposition 2.4. The first part of the theorem is a simple statement of facts from convex analysis, and the strong duality result is based on Lagrange multiplier theory for real valued functions in a reflexive Banach space. In particular, we establish the existence of a solution to Problem (P) (in L_2) under Assumption 2.2 via the direct method (see e.g., [9, thm. 5.51]) and strong duality via the methods of [9, Ch. 9].

Consider a Lagrangian of Problem (P):

$$\begin{aligned}
\mathcal{L}(s, r, \rho, \mu) = & \sum_{j=1}^{M} \int_0^{T^j} \Lambda_j(s_j(t), t) \mathrm{d}t + \sum_{i=1}^{N} \rho_i \Big[C_i - \sum_{j \in \mathcal{A}_i} \int_0^{T_i} r_{ij}(t) \mathrm{d}t \Big] \\
& + \sum_{j=1}^{M} \int_0^{T^j} \mu_j(t) \Big[\sum_{i \in \mathcal{B}_j} r_{ij}(t) - s_j(t) \Big] \mathrm{d}t
\end{aligned} \tag{19}$$

with the problem domain $S = \mathrm{dom}\, \mathcal{L}$ restricted to a convex subset of L_2:

$$S \triangleq \mathcal{P} \times \mathcal{D} \triangleq \{(r, s), (\rho, \mu) \mid r_{ij}(t) \geq 0, \rho_i \geq 0, \mu_j(t) \geq 0\}.$$

It is well known that the primal problem (P) can be posed as finding s, r which attains the infima:

$$\inf_{s, r \in \mathcal{P}} \sup_{\rho, \mu \in \mathcal{D}} \mathcal{L},$$

since the suprema is ∞ whenever the constraints of (P) are not met. Moreover, it is a general theorem (the max-min inequality) that

$$D^\star \triangleq \sup_{\rho,\mu \in \mathcal{D}} \inf_{s,r \in \mathcal{P}} \mathcal{L} \leq \inf_{s,r \in \mathcal{P}} \sup_{\rho,\mu \in \mathcal{D}} \mathcal{L} \triangleq P^\star.$$

We proceed to show that finding ρ, μ to maximize $\inf_{s,r \in \mathcal{P}} \mathcal{L}$ is equivalent to (D). Consider first the r variable, we have:

$$\inf_{r \geq 0} \sum_{i=1}^{N} \sum_{j \in \mathcal{A}_i} \int_0^{T_i} \left(\mu_j(t) - \rho_i \right) r_{ij}(t) \mathrm{d}t,$$

which is $-\infty$ unless $\mu_j(t) \geq \rho_i$ for $i \in \mathcal{B}_j$. Next, examine $\inf_s \mathcal{L}$:

$$\inf_s \sum_{j=0}^{M} \int_0^{T^j} \left[\Lambda_j(s_j(t), t) \mathrm{d}t - \mu_j(t) s_j(t) \right],$$

which can be computed point-wise (it will be clear that the resulting function is in L_2) as

$$\inf_s \left[\Lambda_j(s_j, t) - \mu_j s_j \right] = -\sup_s \left[\mu_j s_j - \Lambda_j(s_j, t) \right]$$
$$= -\Lambda_j^*(\mu_j, t),$$

the negative of the conjugate of $s \mapsto \Lambda_j(s, t)$ (see Proposition 2.2). Combining these calculations together results in the following problem:

$$
\begin{aligned}
\underset{\rho,\mu}{\text{maximize}} \quad & -\sum_{j=1}^{M} \int_0^{T^j} \Lambda_j^*(\mu_j(t), t) \mathrm{d}t + \sum_{i=1}^{N} \rho_i C_i \\
\text{subject to} \quad & \mu_j(t) \geq \rho_i \ \forall i \in \mathcal{B}_j \cap \mathcal{T}_t \\
& \rho_i \geq 0.
\end{aligned}
\tag{20}
$$

Suppose that ρ_i is optimal for this problem. Then, since Λ^* is monotone increasing, the optimal $\mu_j(t)$ is the smallest possible, which implies we can take $\mu_j(t)$ to be piecewise constant between contract deadlines (any deviation over a positive measure interval increases the cost, and any deviation over a measure zero set doesn't change the cost). Denote by μ_{jk} the value taken by the optimal $\mu_j(t)$ over $[T_{k-1}, T_k]$ for $T_k \leq T^j$ to obtain Problem (D). Indeed, it can be seen that $\mu_{jk} = \max_{i \in \mathcal{B}_j \cap \mathcal{T}_{T_k}} \rho_i$.

Finally, we show the existence of solutions to both (P) and (D), as well as the equality of their values, under the adequate supply assumption 2.2.

Under Assumption 2.2 the $\mathbb{R}$-valued function on $L_2(\mathbb{R})^M \times L_2(\mathbb{R})^{N \times M}$ defined by $(s, r) \mapsto \sup_{\rho,\mu} \mathcal{L}(s, r, \rho, \mu)$ is a proper, convex, and l.s.c. function. That this function is proper follows from the assumption since there is at least one feasible point, the convexity follows since suprema over families of affine functions are convex, and lower semi-continuity follows since this property is preserved through the suprema. Naturally, the restriction of S to the primal variables (s, r) remains closed and convex.

Now, by the basic properties of W_j (Assumption 2.1) $\Lambda_j(s)$ is positively-coercive (i.e. $\Lambda_j(s) \to \infty$ as $s \to \infty$); this property is thus inherited by the function $s \mapsto \sup_{\rho,\mu} \mathcal{L}(s, r, \rho, \mu)$. Indeed, by the construction of $\mathcal{L}$ the function $r \mapsto \sup_{\rho,\mu} \mathcal{L}(s, r, \rho, \mu)$ must be positively-coercive as well since if $\sum_{i \in \mathcal{B}_j} r_{ij}(t) \neq s_j(t)$ then $\sup_{\rho,\mu} \mathcal{L}(s, r, \rho, \mu) = \infty$. Then, keeping in mind that $(s, r) \in S \implies s, r \geq 0$, this is enough to satisfy the growth condition required by [9, thm. 5.51 (c)] and establishes the existence of a solution to $\inf_{s,r} \sup_{\rho,\mu} \mathcal{L}$ and therefore to (P).

The statement of strong duality now follows from combining [9, thm. 9.8] and [9, thm. 9.13] with Assumption 2.2. $\square$

C.2 Additional Proofs

Proposition 2.1 (Hierarchy of α-concavity). *For $0 \le \alpha < \beta$, if W is α-concave, then it is also β-concave.*

Proof of Proposition 2.1. First we check that $\ell_\beta \circ \ell_\alpha^{-1}(x)$ is both monotone increasing and concave. This follows if the first derivative

$$\frac{\mathrm{d}}{\mathrm{d}x} \ell_\beta \circ \ell_\alpha^{-1}(x) = \frac{\ell_\beta' \circ \ell_\alpha^{-1}(x)}{\ell_\beta' \circ \ell_\alpha^{-1}(x)},$$

is positive and monotone non-increasing; which, since ℓ_α^{-1} is itself monotone non-decreasing, follows if

$$\frac{\ell_\beta'(x)}{\ell_\alpha'(x)}$$

is both positive and monotone decreasing. Since $\ell_\alpha'(x) = x^{-\alpha}$ this function is $x^{\alpha - \beta}$, which is positive on the domain $x > 0$ and decreasing if $\alpha < \beta$.

Now, we check concavity of $\ell_\beta \circ W$ directly from the definition, for $t \in (0, 1)$:

$$
\begin{aligned}
\ell_\beta \circ W(tx + (1-t)y) &= \ell_\beta \circ \ell_\alpha^{-1} \circ \ell_\alpha \circ W(tx + (1-t)y) \\
&\overset{(a)}{\ge} \ell_\beta \circ \ell_\alpha^{-1}\big(t\ell_\alpha \circ W(x) + (1-t)\ell_\alpha \circ W(y)\big) \\
&\overset{(b)}{\ge} t\ell_\beta \circ \ell_\alpha \circ \ell_\alpha^{-1} \circ W(x) + (1-t)\ell_\beta \circ \ell_\alpha^{-1} \circ \ell_\alpha \circ W(y) \\
&= t\ell_\beta \circ W(x) + (1-t)\ell_\beta \circ W(y),
\end{aligned}
$$

where (a) follows by the assumed concavity of $\ell_\alpha \circ W$ and the monotonicity of $\ell_\beta \circ \ell_\alpha^{-1}$ while (b) from the concavity of $\ell_\beta \circ \ell_\alpha^{-1}$. $\qquad\square$

Proposition 2.5 (Connection Between First and Second Price Case). *Under Assumption 2.2, and whenever the acquisition cost curves are convex, the Problem (TP) for a second price auction is equivalent to solving*

$$
\begin{aligned}
\underset{x,\gamma,\rho}{\text{minimize}} \quad & \sum_{j=1}^{M} \int_0^{T^j} x_j(t) W_j(x_j(t), t)\mathrm{d}t - \sum_{i=1}^{N} \rho_i C_i \\
\text{subject to} \quad & \sum_{j \in \mathcal{A}_i} \int_0^{T_i} \gamma_{ij}(t) W_j(x_j(t), t) \ge C_i \\
& \sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) = 1, \gamma_{ij}(t) \ge 0, \rho_i \ge 0 \\
& x_j(t) \ge \rho_i \; \forall i \in \mathcal{B}_j \cap \mathcal{T}_t.
\end{aligned}
\tag{9}
$$

The objective function involves $x_j(t) W_j(x_j(t), t)$, which is the cost associated with a first price auction.

Proof of Proposition 2.5. We first examine the conjugate function for the second price acquisition function. We have

$$
\Lambda_{2nd}(s) \overset{\triangle}{=} \begin{cases} \infty; & s > \bar{s} \\ 0; & s < W(0) \\ \int_{W(0)}^s W^{-1}(y)\mathrm{d}y; & \text{otherwise} \end{cases}.
\tag{21}
$$

We can find the conjugate on a segment of $\mathbb{R}$ by differentiation:

$$
\Lambda_{2nd}^*(\mu) = \sup_s \big[s\mu - \Lambda(s)\big],
$$

leading to the supremum being attained at $s = W(\mu)$ as long as $\mu \in [0, \bar{x}]$ (in which case the derivative is indeed $W^{-1}(s)$). We already know from Proposition 2.2 that when $\mu < 0$ we have $\Lambda^*(\mu) = \infty$ and if $\bar{s} < \infty$ and $\mu > \bar{x}$, then $\Lambda^*(\mu) = \bar{s}\mu - \Lambda(\bar{s})$.

Therefore, since $\Lambda(W(\mu)) = f(\mu)$ by definition, we have

$$\Lambda^*_{2nd}(\mu) = \begin{cases} \infty; & \mu < 0 \\ \mu\bar{s} - \Lambda_{2nd}(\bar{s}); & \mu > \bar{x} \\ \mu W(\mu) - f_{2nd}(\mu); & \text{otherwise} \end{cases}. \tag{22}$$

Given a dual solution μ, ρ of Problem (D), which attains the primal problem's value by strong duality Proposition 2.4, we see from the above calculations that the maximizing argument for Λ^* is just $s = W(\mu)$, therefore it must be the case that the optimal bid is the optimal multiplier $x_j(t) = \mu_{jk}$ on $t \in [T_{k-1}, T_k)$, which recovers the result of Example 2.1 that for the second price auction there exists an optimal piecewise constant bid.

Using this fact, and writing the dual as a minimization problem by negating its objective, we can consider the problem of *simultaneously* solving the dual (D) and the original Problem (TP)

$$\begin{aligned} \underset{x,\gamma,\rho}{\text{minimize}} \quad & \sum_{j=1}^{M} \int_{0}^{T^j} \left[f_j^{2nd}(x_j(t), t) + \Lambda_j^*(x_j(t), t) \right] \mathrm{d}t - \sum_{i=1}^{N} \rho_i C_i \\ \text{subject to} \quad & \sum_{j \in \mathcal{A}_i} \int_{0}^{T_i} \gamma_{ij}(t) W_j(x_j(t), t) \geq C_i \\ & \sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) = 1, \gamma_{ij}(t) \geq 0, \rho_i \geq 0 \\ & x_j(t) \geq \rho_i \,\forall i \in \mathcal{B}_j \cap \mathcal{T}_t, \end{aligned} \tag{23}$$

where we have again applied the fact that there exists a solution such that $x_{ij}(t) = x_j(t)$ for every $i \in \mathcal{B}_j$ and where $\sum_{i \in \mathcal{B}_j} \gamma_{ij}(t) = 1$ (see the proof of 2.3).

If we can show that the optimal dual multipliers satisfy $\mu_{jk} \leq \bar{x}_j$, then using the fact that $x_j(t) = \mu_{jk}$ for $t \in [T_{k-1}, T_k)$ and the expression for Λ^* we have the result, since the f_j^{2nd} terms cancel.

To this end, recognize that since $\mu_{jk} = \max_{i \in \mathcal{B}_j \cap \mathcal{T}_{T_k}} \rho_i$, we must have $\mu_{jk} \leq \bar{x}_j$ unless there is some $\rho_i > \bar{x}_j$. We suppose by way of contradiction that $\rho_i > \bar{x}_j$. Consider the portion of the objective of (D) related to contract i (where $\Lambda = \Lambda_{2nd}$):

$$\begin{aligned} -\sum_{j \in \mathcal{A}_i} \sum_{k: T_k \leq T_i} \int_{T_{k-1}}^{T_k} \Lambda_j^*(\mu_{jk}, t) \mathrm{d}t + \rho_i C_i & \overset{(a)}{\leq} \rho_i C_i - \sum_{j \in \mathcal{A}_i} \int_{0}^{T_i} \Lambda_j^*(\rho_i, t) \mathrm{d}t \\ & \overset{(b)}{=} \rho_i C_i - \sum_{j \in \mathcal{A}_i} \int_{0}^{T_i} (\rho_i \bar{s}_j(t) - \Lambda_j(\bar{s}_j(t), t)) \mathrm{d}t \\ & = \sum_{j \in \mathcal{A}_i} \int_{0}^{T_i} \Lambda_j(\bar{s}_j, t) \mathrm{d}t + \rho_i \left[C_i - \int_{0}^{T_i} \bar{s}_j(t) \mathrm{d}t \right] \\ & \overset{(c)}{<} \sum_{j \in \mathcal{A}_i} \int_{0}^{T_i} f_j(\bar{x}_j(t), t) \mathrm{d}t \end{aligned}$$

where (a) follows from $\mu_{jk} \geq \rho_i$ since Λ^* is non-decreasing, (b) from Equation (22) since $\rho_i > \bar{x}_j$, and (c) from $\rho_i \geq 0$ combined with Assumption 2.2 which ensures the bracketed term is negative. The final expression is the cost attained by bidding $\bar{x}_j(t)$ so by strong duality $\rho_i > \bar{x}_j$ must not have been optimal (recall that in the dual we are maximizing). Therefore, $\mu_{jk} \leq \bar{x}_j$, which completes the proof. $\qquad \square$

Proposition 3.1 (Finite Primal Problem). *The finite optimization problem over the variables $r_{ij}[k], s_{ij}[k]$ defined by*

$$\underset{s,r}{\text{minimize}} \quad \sum_{j=1}^{M} \sum_{k:\widetilde{T}_k \leq T^j} \overline{\Lambda}_{jk}(s_j[k])$$

$$\text{subject to} \quad \sum_{j\in\mathcal{A}_i} \sum_{k:\widetilde{T}_k \leq T_i} \Delta_k r_{ij}[k] \geq C_i \qquad\qquad (P_K)$$

$$\sum_{i\in\mathcal{B}_j\cap\mathcal{T}_{T_k}} r_{ij}[k] = s_j[k]$$

$$r_{ij}[k] \geq 0,$$

is a finite approximation of (P)*: any solution* $(s_j[k], r_{ij}[k])$ *of* (P_K) *is feasible for Problem* (P) *(as a piecewise constant function of* t*).*

Let $g_j^{-1}(\mu, t)$ *denote a continuous selection from Equation (8) of Example 2.1 (e.g., choose the right derivative). If each cost function* $f_j(x, t) = xW_j(x, t)$ *is* L_f*-Lipschitz in* x *(uniformly in* t *a.e.) and* $g_j^{-1}(x, t)$ *is* Γ_g*-Lipschitz in* t *a.e. (uniformly in* x*), then the cost difference* ϵ_K *of the optimal solution to* (P) *and the feasible piecewise constant approximation obtained from Problem* (P_K) *is bounded by*

$$\epsilon_K \leq \frac{M}{4} L_f \Gamma_g \sum_{k=1}^{K} \Delta_k^2.$$

If for each j *the function* $W_j(x, t)$ *(and hence* $\Lambda_j(s, t)$*) is almost everywhere twice continuously differentiable in* t*, the integral approximation error* δ_K *between the objective of* (P) *and* (P_K) *is bounded by* $\delta_K = \mathcal{O}(\sum_{k=1}^{K} \Delta_k^3)$*.*

In particular, if we have $\Delta_k = T/K$ *(valid as long as* $\widetilde{T}_k$ *contains the contract deadlines), then we have the bounds*

$$\epsilon_K \leq \frac{M L_f \Gamma_g T^2}{4K}$$
$$\delta_K = \mathcal{O}(1/K^2).$$

The total discretization error between the value of (P_K) *and the value of* (P) *is therefore* $\epsilon_K + \delta_K = \mathcal{O}(1/K)$*.*

Proof of Proposition 3.1. Let $(s^\star[k], r^\star[k])$ be an optimal solution to Problem (P_K) and $(s^\star(t), r^\star(t))$ an optimal solution to (P) with corresponding (piecewise constant) dual multipliers $\mu_j^\star(t)$. It is evident that $(s^\star[k], r^\star[k])$ is feasible for the continuous time problem (P) by construction since the integral of a piecewise constant function is e.g., $\Delta_k r_{ij}[k]$.

Bounding the integral approximation error is a simple application of the well known error bound for the trapezoidal rule (see e.g. [31, ch. 7]), i.e., for any point s_0 there is some $\hat{t} \in (\widetilde{T}_{k-1}, \widetilde{T}_k)$ such that

$$\Big| \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} \Lambda_j(s_0, t) \mathrm{d}t \Big| \leq \frac{\Delta_k^3}{24} \frac{\partial^2 \Lambda_j(s_0, t)}{\partial t^2} \Big|_{t=\hat{t}}.$$

Now the integral approximation error, i.e., the error between the objective of (P_K) and (P) evaluated on the piecewise constant solution $s^\star[k]$ of (P_K) is

$$\Big| \sum_{j=1}^{M} \sum_{k:\widetilde{T}_k \leq T^j} \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} \Lambda_j(s^\star[k], t) \mathrm{d}t - \sum_{j=1}^{M} \sum_{k:\widetilde{T}_k \leq T^j} \overline{\Lambda}_{jk}(s^\star(\widetilde{T}_k)) \Big| \leq \mathcal{O}\Big(M \sum_{k=1}^{K} \Delta_k^3 \Big).$$

To bound the error of using the solution to (P_K) as a feasible point for (P), we would like to find a point in time where attained supply is maximized:

$$\tau_k \in \Big\{ t \in [\widetilde{T}_{k-1}, \widetilde{T}_k) \mid s^\star(\tau_k) \geq s^\star(u) \; \forall u \in [\widetilde{T}_{k-1}, \widetilde{T}_k] \Big\},$$

31

though this set may be empty since $s^\star(t)$ can be discontinuous at $\widetilde{T}_k$. However, since we are here assuming smoothness of the supply rate curves in t, the total supply attained in an ϵ-interval of time must converge to 0 as $\epsilon \to 0$ and therefore there must exist points $\tau_k \in [\widetilde{T}_{k-1}, \widetilde{T}_k)$ such that the piecewise constant function $k \mapsto s^\star(\tau_k)$ is feasible for (P_K), since $s^\star(t)$ is feasible for (P). We therefore have the inequalities:

$$\int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} \Lambda_j(s_j^\star(t), t) \mathrm{d}t \le \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} \Lambda_j(s_j^\star[k], t) \mathrm{d}t \le \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} \Lambda_j(s_j^\star(\tau_k), t) \mathrm{d}t.$$

We can bound the absolute difference between the outer terms as follows:

$$
\begin{aligned}
\int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} |\Lambda_j(s_j^\star(t), t) - \Lambda_j(s_j^\star(\tau_k), t)| \mathrm{d}t &\overset{(a)}{=} \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} |f_j(g_j^{-1}(\mu_j^\star(t), t), t) - f_j(g_j^{-1}(\mu_j^\star(t), \tau_k), t)| \mathrm{d}t \\
&\overset{(b)}{\le} L_f \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} |g_j^{-1}(\mu_j^\star(t), t) - g_j^{-1}(\mu_j^\star(t), \tau_k)| \mathrm{d}t \\
&\overset{(c)}{\le} L_f \Gamma_g \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} |t - \tau_k| \mathrm{d}t \\
&\le \frac{1}{4} L_f \Gamma_g \Delta_k^2,
\end{aligned}
$$

where (a) follows by definition of Λ_j and g^{-1}, (b) follows by the Lipschitz continuity of $f(x, t)$ w.r.t. x; and (c) follows by the Lipschitz continuity of $g^{-1}(x, t)$ w.r.t. t.

Summing these individual error terms we obtain an overall error bound

$$\sum_{j=1}^{M} \left| \int_0^T \Lambda(s^\star(t), t) \mathrm{d}t - \sum_{k=1}^{K} \Delta_k \overline{\Lambda}_k(s^\star[k]) \right| \le \frac{M}{4} L_f \Gamma_g \sum_{k=1}^{K} \Delta_k^2.$$

The total error then arises from summation

$$
\begin{aligned}
\left| \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} \Lambda_j(s^\star(t)) \mathrm{d}t - \overline{\Lambda}_j(s^\star[k]) \right| &\le \left| \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} \Lambda_j(s^\star(t)) \mathrm{d}t - \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} \Lambda_j(s^\star(\widetilde{T}_k)) \right| \\
&\quad + \left| \int_{\widetilde{T}_{k-1}}^{\widetilde{T}_k} \Lambda_j(s^\star(\widetilde{T}_k k)) - \overline{\Lambda}_j(s^\star(\widetilde{T}_k)) \right|,
\end{aligned}
$$

and results in a total error $\epsilon_K + \delta_K = \mathcal{O}(1/K)$. $\qquad\square$