

Rate of Price Discovery in Iterative Combinatorial Auctions

JACOB ABERNETHY, University of Michigan
SÉBASTIEN LAHAIE, Microsoft Research
MATUS TELGARSKY, University of Michigan

We study a class of iterative combinatorial auctions which can be viewed as subgradient descent methods for the problem of pricing bundles to balance supply and demand. We provide concrete convergence rates for auctions in this class, bounding the number of auction rounds needed to reach clearing prices. Our analysis allows for a variety of pricing schemes, including item, bundle, and polynomial pricing, and the respective convergence rates confirm that more expressive pricing schemes come at the cost of slower convergence. We consider two models of bidder behavior. In the first model, bidders behave stochastically according to a random utility model, which includes standard best-response bidding as a special case. In the second model, bidders can behave arbitrarily (even adversarially), and meaningful convergence relies on properly designed activity rules.

1. INTRODUCTION.

Combinatorial auctions are used to sell multiple distinct items at once in situations where items may be substitutes or complements. Because of their generality as a resource allocation mechanism, combinatorial auctions have been proposed for a variety of domains including the allocation of wireless spectrum, airport landing slots, and real estate [Cramton et al. 2006]. In an *iterative* combinatorial auction, bidders place bids on bundles of items in response to prices updated by the auctioneer, and this process repeats until bidding reaches quiescence. Iterative designs are particularly attractive for combinatorial auctions because the sheer size of the bundle space makes it impractical to report valuation information in a single shot.

At the core of an iterative auction design is the choice of pricing scheme, because prices drive the information revelation process. The current space of auctions offers two extremes: linear (item) pricing, where the price of a bundle is the sum of its item prices, and bundle pricing, where each bundle is explicitly priced. The two schemes have different advantages. Linear prices provide information about final costs even for bundles not explicitly bid on, leading to fewer rounds of bidding—in this sense, they provide effective ‘price discovery’. However, in the presence of complementarities linear prices cannot effectively balance supply and demand, leading to inefficiencies in the final allocation. Bundle prices, on the other hand, can always clear the market and support an efficient allocation, but provide limited price discovery. As a result, bundle-price auctions typically require far more rounds in practice to reach termination, as confirmed empirically in both simulation studies and lab experiments [Scheffel et al. 2011; Schneider et al. 2010].

In this paper we consider the design of iterative auctions from an algorithmic perspective, leading to a formal study of the relationship between price structure and price discovery. Prior work has shown that iterative multi-item auctions generally fall under two design paradigms: they can be viewed as either primal-dual or subgradient algorithms to solve the dual problems of allocation and pricing, once these are formulated as linear programming problems [Bikhchandani et al. 2002]. In this work we consider auctions based on the subgradient method. We obtain general linear and quadratic programming formulations of the pricing problem solved by the auction, which allows

Author addresses: J. Abernethy, Department of Computer Science, University of Michigan; email: jabernet@umich.edu; S. Lahaie, Microsoft Research, New York City; email: slahaie@microsoft.com; M. Telgarsky, Department of Computer Science, University of Michigan; email: mtelgars@umich.edu

us to study various price structures under a single framework. We provide a general specification of subgradient auctions which subsumes several prior auctions using linear or bundle prices, and also leads to novel auction designs using polynomial prices in between these two schemes. Polynomial prices extend linear prices by assigning coefficients to combinations of items, and can allow the auctioneer to strike a more careful balance between market clearing and price discovery.

Our main focus is the convergence rate of iterative auctions, and in particular its dependence on the pricing scheme used. We consider two agent models. In the first model, agents bid on their most preferred bundles at the current prices, but their value estimates for bundles are subject to stochastic errors at each round. This translates into stochastic behavior where bids may be placed on lesser-preferred bundles by chance. Our convergence analysis confirms that subgradient auctions reach clearing prices even under this kind of imperfect bidding behavior, and bounds the rates for linear, bundle, and polynomial pricing. In the second model, agents can place bids arbitrarily or even adversarially in each round upon seeing the current prices. While it is still possible to prove certain technical convergence statements under such a model, it unsurprisingly admits behavior that would be unreasonable and disallowed in practice. We therefore show how the convergence statements become more meaningful when imposing revealed-preference activity rules (constraints on bids over rounds) that have been proposed in practice [Ausubel and Baranov 2014b].¹

At the core of our analysis are tools and concepts from statistical learning theory, online (sequential) learning, and convex analysis. The auction is cast as an iterative procedure whose goal is to optimize an objective function over prices, in analogy to fitting a learning model to minimize prediction loss. The tension between market clearing and price discovery mirrors the bias-variance trade-off familiar from statistics. Linear prices offer ‘low variance’, which makes them informative about final costs even after a limited number of rounds. Bundle prices offer ‘low bias’, which is needed to flexibly price bundles to balance demand and supply. (The notions of bias and variance here are analogies, because there is no randomness in the auction mechanism.) We emphasize the price discovery aspects throughout the paper. While our analysis does provide structural insights into the market clearing properties of various pricing schemes, a detailed treatment of this question requires one to consider restricted valuation classes, which we defer to a separate study.

The remainder of the paper is organized as follows. Section 2 introduces the elements of our model and describes our approach to price representation. Section 3 formulates the core problems of allocation and pricing, and develops the duality relationship between the two. Section 4 specifies the class of auctions based on the subgradient method. The main results of the paper appear in Sections 5 and 6, providing convergence rates for the auctions under the stochastic and arbitrary bidding models, respectively. Section 7 concludes. Detailed proofs of all results are deferred to the appendices. Sketches of the proofs are given in the main text.

Related work. There is an extensive literature on the design of multi-item auctions [e.g. Milgrom 2004]. Our paper relates to the narrower literature on the algorithmic properties of iterative auctions. The allocation problem at the basis of combinatorial auctions was first formulated in a linear programming (LP) framework

¹Under both the stochastic model and the adversarial model with activity rules, agents bid consistently with some underlying valuation for the bundles, but it need not correspond to their true valuation. Agents may benefit from strategic bidding. This has no bearing on the auction’s worst-case convergence rate, so we set aside incentive concerns in this work. As with other iterative auctions used in practice, supplementary pricing stages such as core-selecting or VCG pricing may be used to achieve reasonable outcomes in equilibrium [Day and Cramton 2012].

by Bikhchandani and Ostroy [2002], who provide formulations leading to linear and bundle prices, both anonymous and personalized. The formulation used in this paper is equivalent to the one used by Lahaie [2011] for his primal-dual auction, and subsumes the LPs of Bikhchandani and Ostroy [2002]. The work of Lahaie [2009, 2011] introduced the use of polynomial prices for combinatorial auctions.

Drawing on the LP viewpoint, Bikhchandani et al. [2002] and de Vries et al. [2007] categorized existing auctions as either primal-dual or subgradient algorithms. Subgradient auctions include the uniform price auction for homogeneous goods, and the well-known auctions of Crawford and Knoer [1981], Kelso and Crawford [1982], Parkes [1999], and Ausubel and Milgrom [2002] for heterogeneous goods. All of these auctions are part of the class studied here (for a suitable choice of step-size policy).

An important aspect of our auction class is that we do not enforce price monotonicity: the price of a bundle may ascend and descend during the course of the auction. While there is precedent for non-monotone price paths [e.g. Ausubel 2006], most auction designs favor ascending prices to rule out certain gaming behaviors on the part of bidders that aim to delay termination. We show how activity rules can mitigate such concerns and restore convergence. Ausubel and Baranov [2014b] discuss the practical aspects of combinatorial auction design, while Ausubel and Cramton [2011] consider activity rules. Harsha et al. [2010] provide a detailed treatment of revealed preference activity rules similar to the rule considered in this paper.

Many of the core ideas used herein draw from convex analysis [Borwein and Lewis 2010; Hiriart-Urruty and Lemaréchal 2001; Rockafellar 1970] and the mathematical results that explore sequential optimization procedures such as subgradient descent, or more broadly *mirror descent* [Beck and Teboulle 2003]. Many of these methods have recently been employed within the machine learning community, where prediction and decision problems are frequently viewed as an *online convex optimization game* [Zinkevich 2003]. In this model, a learner is repeatedly asked to select decisions from a convex set, and on each round the feedback on this decision arrives in the form of a convex loss function, potentially selected by an adversary. In the context of iterative auctions, the decision is the price vector at each round, and the resulting bids provide feedback on the ‘loss’ incurred. Online convex optimization games have been thoroughly explored in recent years and one can find a number of comprehensive surveys, including the work of Cesa-Bianchi and Lugosi [2006], Hazan [2012], and Shalev-Shwartz [2011].

2. THE FORMAL MODEL.

We consider a model with n agents (buyers) and a single seller holding m distinct items. At a high level, the purpose of the auction is to allocate the m items to the n agents according to how the agents value the items. We use the notation $[n] = \{1, 2, \dots, n\}$ to denote an index set; thus $[n]$ indexes the set of agents and $[m]$ indexes the set of items. To formulate our model, we first treat the items as divisible; we will later explicitly impose indivisibility (i.e, integrality) requirements. Agents have preferences over various bundles of items, where a bundle is a subset of the items. Let X denote the set of all bundles the agents would be interested in acquiring, and let $\ell = |X|$. In general, this could be all possible bundles, in which case $\ell = 2^m$, but in some applications it may be substantially smaller.²

Notation. We use the following convention to index vectors and matrices throughout the paper. For a vector $a \in \mathbf{R}^\ell$ indexed by the finite set of bundles, we use the functional

²For instance, the literature often considers *single-minded* agents. Such agents only derive positive value if they acquire a specific, designated bundle, and the marginal value of all other items is zero. Under single-minded agents we have $\ell \leq n$.

notation $a(x)$ for the component corresponding to $x \in X$. For a vector $a \in \mathbf{R}^n$ indexed by agents, we use the usual subscript notation a_i for $i \in [n]$. For a vector $a \in \mathbf{R}^{n\ell}$ indexed by both, we write $a_i(x)$ to refer to a component. The convention extends to matrices. For a matrix $A \in \mathbf{R}^{n\ell \times d}$, we write $A_i(x)$ for the d -dimensional row associated with agent i and bundle x .

The primitives of our model will be cast as elements of vector spaces. A bundle assigned to an agent i is represented as a vector from the agent's *consumption set*

$$H_i = \text{conv} \left\{ q_i \in \{0, 1\}^\ell : \sum_{x \in X} q_i(x) \leq 1 \right\}, \quad (1)$$

where ‘conv’ denotes the convex hull operator. The consumption set is a polytope whose extreme points are in one-to-one correspondence with bundles; an extreme point q_i is a binary vector corresponding to the unique $x \in X$ for which $q_i(x) = 1$, or the empty bundle if q_i is the origin. Fractional vectors represent bundles with fractional quantities of items. The agents' consumption sets are identical, but this is not important for our results. We use $H = H_1 \times \dots \times H_n$ to denote the agents' joint consumption set.

An *allocation* is represented by a vector $q = (q_1, \dots, q_n) \in \mathbf{R}^{n\ell}$, where the subvector $q_i \in \mathbf{R}^\ell$ encodes the bundle that agent i receives. An allocation is feasible if no more than one unit of each item is supplied to the agents in total. More formally, the set of feasible allocations is captured by the seller's *production set*

$$F = \text{conv} \left\{ q \in \{0, 1\}^{n\ell} : \sum_{i \in [n]} \sum_{x \in X, x \ni j} q_i(x) \leq 1 \ (j \in [m]), \sum_{x \in X} q_i(x) \leq 1 \ (i \in [n]) \right\}. \quad (2)$$

The production set is a polytope whose extreme points are in one-to-one correspondence with lists of bundles (x_1, \dots, x_n) such that each item appears in at most one bundle and each agent receives no more than one bundle. They therefore correspond to feasible allocations of indivisible items. Note that, by definition, if $q \in F$ then $q_i \in H_i$ for each agent i , so that $F \subset H$.

Each agent i has a *valuation* $v_i \in \mathbf{R}^\ell$ which records the agent's willingness to pay for each bundle in a common unit of currency. We assume that each agent's value for the empty bundle is 0. A *valuation profile* is a vector of agent valuations $v = (v_1, \dots, v_n) \in \mathbf{R}^{n\ell}$. Given a valuation profile v , an allocation \bar{q} is *efficient* (in the economic sense) if it maximizes the total value to the agents among all feasible allocations:

$$\bar{q} \in \arg \max_{q \in F} v^\top q. \quad (3)$$

Like valuation profiles, *prices* are described by a vector $p \in \mathbf{R}^{n\ell}$ where $p_i(x)$ is the charge for bundle x to agent i . As defined the prices may be *personalized*, in the sense that two different agents may see different prices for identical bundles. We assume that agents have *quasi-linear utility*: given prices p , agent i 's utility for bundle $x \in X$ is $v_i(x) - p_i(x)$. Equivalently, if $q_i \in H_i$ is the vector corresponding to bundle x , the utility is $v_i^\top q_i - p_i^\top q_i$. If the auction charges prices p and allocates according to $q \in F$ then the *revenue* totals $p^\top q$. The seller has zero value for the items and only derives utility from the revenue collected.

Price representation. We have so far introduced prices as vectors from $\mathbf{R}^{n\ell}$ which explicitly list the price of each bundle to each agent. The purpose of this paper is to analyze the impact of different pricing schemes (e.g., linear or bundle) on the operation of the auction, especially its rate of convergence. To impose further structure on prices and restrict them to a lower-dimensional subspace, we use an indirect approach

that first defines an alternative vector space representation for the bundles. This approach, explored in depth in [Lahaie 2009, 2011], draws very much from the class of *kernel methods* used in machine learning. In the context of prediction and estimation, the intuition behind kernel methods is that they map the data to a (potentially high dimensional) space where the function to be learned is *linear* in the space. The same idea applies in the context of our auction design, where our goal is to produce a pricing function that is linear under a particular representation of the bundles.

We introduce a *representation matrix* $G \in \mathbf{R}^{n\ell \times d}$, where $m \leq d \leq n\ell$. The row $G_i(x)$ provides a d -dimensional encoding of bundle $x \in X$, which can also depend on the agent i in general. The interpretation of the encodings is that they define the “features” of the bundles that are priced in the auction. We take prices to be linear functions of bundle encodings, so that prices can be represented in \mathbf{R}^d . That is, the price vector $p \in \mathbf{R}^{n\ell}$ can be written as $p = Gw$ for some *price parameter* vector $w \in \mathbf{R}^d$; we will sometimes drop p and refer to the prices as Gw . As a convention, the empty bundle is always encoded as the origin in \mathbf{R}^d , which means that its price is normalized to 0. To make these ideas more concrete, let us consider several examples of representations, each leading to different pricing schemes.

Linear. There is a feature for each item ($d = m$). A bundle is encoded using its standard 0-1 indicator vector representation. For example, with items a, b, c , the representations of bundles $\{a, b\}$ and $\{a, b, c\}$ are respectively

$$\begin{aligned} \{a, b\} &\mapsto \begin{bmatrix} a & b & c \\ 1 & 1 & 0 \end{bmatrix}, \\ \{a, b, c\} &\mapsto \begin{bmatrix} 1 & 1 & 1 \end{bmatrix}. \end{aligned}$$

This leads to simple item pricing: the price of a bundle is the sum of the prices associated with each of its items. In the sequel, the notation $G_{p(1)}$ will refer to linear pricing matrices (polynomials of degree 1).

Bundle. At the other extreme, we can have a feature for each relevant non-empty bundle ($d = \ell$). A non-empty bundle is encoded by the unit vector which has a 1 in the component corresponding to the bundle. For example, with items a, b, c , and assuming X consists of all bundles, the representations of bundles $\{a, b\}$ and $\{a, b, c\}$ are respectively

$$\begin{aligned} \{a, b\} &\mapsto \begin{bmatrix} a & b & c & ab & ac & bc & abc \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}, \\ \{a, b, c\} &\mapsto \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \end{aligned}$$

Here each bundle is explicitly priced and there may be no relationship whatsoever between the prices of different bundles. In the sequel, G_{id} will refer to bundle pricing matrices.

Polynomial. As a generalization of the linear representation, we can have a feature for subsets of items up to size r , for $1 \leq r \leq m$. A component is 1 if the bundle contains all the items in the associated subset, and 0 otherwise. For example, with items a, b, c , the representations of bundles $\{a, b\}$ and $\{a, b, c\}$ using $r = 2$ are respectively

$$\begin{aligned} \{a, b\} &\mapsto \begin{bmatrix} a & b & c & ab & ac & bc \\ 1 & 1 & 0 & 1 & 0 & 0 \end{bmatrix}, \\ \{a, b, c\} &\mapsto \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}. \end{aligned}$$

Prices therefore take the form of multi-variate polynomials³ of degree r . Note that price components can be positive or negative, so polynomial prices of degree $r \geq 2$ may be super- or sub-additive on various bundles. In the sequel, $G_{p(r)}$ will denote polynomial pricing matrices of degree r .

Note that while the bundle representation is fully expressive, it is conceptually distinct from the linear and polynomial representations and should not be construed as a generalization of either. We are not aware of any iterative auctions in the literature using price structures other than the ones listed above. However, the formalism can also accommodate intricate pricing schemes such as attribute pricing (e.g., square footage for real estate, or population density for wireless spectrum). The encodings cannot be entirely arbitrary: they must retain enough information about the items contained in the bundles to allow one to verify, given a vector of n bundle encodings, whether it forms a feasible allocation. The above encodings all have this property—see [Lahaie 2011] for a more detailed discussion.

The representations given above (linear, bundle, and polynomial) are all anonymous, in the sense that $G_i(x) = G_j(x)$ for all $i, j \in [n]$, leading to anonymous prices. To personalize prices, we can include features that depend on the identity of the agent. For instance, one can introduce an agent intercept to any of the encodings above, leading to a feature space of dimension $d+n$. To obtain entirely separate prices for each agent, one can take n copies of the original feature space and map a bundle into the agent’s copy; the resulting dimension is nd . Our results can be adapted to such schemes by extending the dimension of the feature space accordingly.

3. ALLOCATION AND PRICING.

It has been understood since the literature on stability of equilibria that price adjustment processes minimize a convex potential involving indirect utilities [Arrow and Hahn 1971; Varian 1981]. In the context of combinatorial auctions, this potential arises as the dual of the efficient allocation problem [Bikhchandani et al. 2002; Bikhchandani and Ostroy 2002]. We give a full development of this duality for our setup. This will serve to clarify why iterative auctions can be analyzed as subgradient methods, and also provides bounds on price magnitudes, important for later analysis.

Consider the agents’ *aggregate* indirect utility function and demand correspondence, defined respectively as

$$u(p; v) = \max_{q \in H} v^\top q - p^\top q, \quad (4)$$

$$U(p; v) = \arg \max_{q \in H} v^\top q - p^\top q, \quad (5)$$

for personalized bundle prices $p \in \mathbf{R}^{n\ell}$. We will often suppress parameter v when clear from context. As H is a product set, an optimal q in (4–5) decomposes into $q = (q_1, \dots, q_n)$ where q_i maximizes agent i ’s utility $v_i^\top q_i - p_i^\top q_i$ individually over H_i . The indirect utility u therefore aggregates the agents’ maximal utilities over bundles, given prices p . Similarly, correspondence U maps to vectors of utility-maximizing bundles, listing one bundle for each agent, which may overlap in general. We also have the

³To see this, note that the price function in this example can be written as the quadratic polynomial $w_a x_a + w_b x_b + w_c x_c + w_{ab} x_a x_b + w_{ac} x_a x_c + w_{bc} x_b x_c$, where x_a, x_b, x_c are binary indicator variables for the items contained in the bundle, and w is the price parameter vector.

seller's indirect utility function and supply correspondence, defined respectively as

$$s(p) = \max_{q \in F} p^\top q, \quad (6)$$

$$S(p) = \arg \max_{q \in F} p^\top q. \quad (7)$$

We say that prices \bar{p} are *market clearing* if $U(\bar{p}) \cap S(\bar{p}) \neq \emptyset$, and for any allocation $\bar{q} \in U(\bar{p}) \cap S(\bar{p})$ we say that prices \bar{p} *support* \bar{q} . More explicitly, this means that: 1) $\bar{q} \in F$, so it represents a feasible allocation; 2) the assigned bundles according to \bar{q} maximize the agents' utilities at prices \bar{p} ; 3) allocation \bar{q} maximizes the seller's revenue at prices \bar{p} . Since the agents and seller would each willingly select allocation \bar{q} when faced with prices \bar{p} , the prices are market clearing in this sense. If there exists prices \bar{p} supporting an allocation \bar{q} , then the allocation is efficient, because for any $q \in F$ we have

$$v^\top \bar{q} = (v^\top \bar{q} - \bar{p}^\top \bar{q}) + \bar{p}^\top \bar{q} \geq (v^\top q - \bar{p}^\top q) + \bar{p}^\top q = v^\top q, \quad (8)$$

where the inequality holds because $\bar{q} \in U(\bar{p})$ and $\bar{q} \in S(\bar{p})$. An iterative auction proceeds by updating a provisional allocation and prices given agent bids in each round, until the prices support the allocation. According to derivation (8), this provides a certificate that an efficient allocation has been reached.

The convex dual to the problem of computing an efficient allocation is the problem of finding market clearing prices. As the minimization problem is on the price parameter, we will refer to the price optimization as the *primal* objective and the allocation optimization as the *dual* objective.⁴ In what follows, 'ker' refers to the kernel (i.e., nullspace) of a matrix. Also, given any two subsets of Euclidean space A, B , the sum $A + B$ is defined in the natural way: $A + B = \{x + y : x \in A, y \in B\}$.

THEOREM 3.1. *Consider the optimization problems*

$$\inf \{u(Gw) + s(Gw) : w \in \mathbf{R}^d\}, \quad (9)$$

$$\sup \{v^\top q : q \in H \cap (\ker(G^\top) + F)\}. \quad (10)$$

Then the primal value in (9) equals the dual value in (10), and both primal and dual optima are attained. Moreover, allocation \bar{w} and price parameter \bar{q} are optimal primal and dual solutions if

$$\bar{q} \in U(G\bar{w}) \cap \left(\ker(G^\top) + S(G\bar{w}) \right). \quad (11)$$

To understand the dual formulation (10), which captures the allocation problem, note first that the objective is the efficiency $v^\top q$ of allocation q . The first part of the constraint is simply $q \in H$, meaning agents are allocated bundles from their consumption sets. This is combined with the constraint $q \in \ker(G^\top) + F$. If the latter set were simply F , then the feasible set would reduce to $H \cap F = F$, which is simply the convex hull of set of feasible allocations. Recall that the role of G is to constrain the possibilities for the pricing space. Dually, restricting the dimension of the rows in G expands $\ker(G^\top)$, and $\ker(G^\top) + F$ becomes a relaxation of F . This discussion leads to the following result.

COROLLARY 3.2. *If G has full row rank, then there is an integer optimal solution \bar{q} to the dual objective (10), and for any primal optimal solution \bar{w} , prices $\bar{p} = G\bar{w}$ support allocation \bar{q} .*

⁴This is consistent with the usual convention in convex analysis, but unfortunately conflicts with the convention in the auctions literature, where the primal is typically the allocation problem and the dual is the pricing problem. We adopt the former convention because it becomes much simpler to directly apply convex analysis results in the proofs.

For instance, G has full row rank if it encodes personalized bundle prices, or personalized polynomial prices of degree m , which shows that such prices can support efficient integer allocations (i.e., allocations of indivisible items). The former case was first proved by Bikhchandani and Ostroy [2002]. If the set of relevant bundles X is restricted, lower-dimensional prices may suffice. The condition given in Corollary 3.2 is sufficient but not necessary, so in practice lower-dimensional prices may still be able to clear the market even when X is large. In order to interpret the results in the sequel, one can assume that G is sufficiently expressive to ensure that (10) has an integer optimal solution; if this is not the case, our results still meaningfully bound convergence rates under divisible items.

Let us now consider the primal more closely. We write $\mathcal{D} : \mathbf{R}^d \rightarrow \mathbf{R}$ to represent the primal objective function over price parameters w , parametrized by the valuation vector v :

$$\mathcal{D}(w; v) = u(Gw; v) + s(Gw).$$

Recall that the prices are $p = Gw$. The first term is aggregate indirect utility, namely the maximum surplus that agents can achieve by each selecting their preferred bundle at prices p (again, the bundles may overlap). Prices should be set high to minimize this term. The second term is the seller's indirect utility, namely the maximum revenue possible over all feasible allocations, under prices p . Prices should be set low to minimize this term. Thus the two terms lead prices to strike a balance between demand and supply. Looking ahead to the iterative auction of Section 4, the auction can be construed as a subgradient method to optimize the primal objective (9), obtaining subgradient information from the agents' bids, which lie in the dual space associated with objective (10).

A difficulty with formulations (9) and (10) is that the optimal prices in the primal may not be unique, as both primal and dual are linear programs. As a result it is not possible analyze convergence of the actual prices, only convergence in objective value. To obtain a strictly convex objective and unique solution, we can introduce a regularization term $\|w\|_2^2/2$ with weight $\lambda > 0$ into the primal objective. The duality result generalizes as follows.

THEOREM 3.3. *Given weight $\lambda > 0$, consider the optimization problems*

$$\inf \left\{ u(Gw) + s(Gw) + \frac{\lambda}{2} \|w\|_2^2 : w \in \mathbf{R}^d \right\}, \quad (12)$$

$$\sup \left\{ v^\top q - \frac{1}{2\lambda} \|G^\top(q - q')\|_2^2 : q \in H, q' \in F \right\}. \quad (13)$$

Then the primal value in (12) equals the dual value in (13), and both the primal and dual optima are attained. Furthermore, the primal optimum \bar{w} is unique. Price parameter \bar{w} and allocations (\bar{q}, \bar{q}') are optimal primal and dual solutions if

$$\bar{q} \in U(G\bar{w}), \quad \bar{q}' \in S(G\bar{w}), \quad G^\top(\bar{q} - \bar{q}') = \lambda\bar{w}. \quad (14)$$

The original duality relation derived in Theorem 3.1 can be viewed as the limiting case of Theorem 3.3 when we set $\lambda = 0$. To make this more transparent, first observe that the constraint in the dual objective (10) can be re-written as

$$G^\top q = G^\top q' \quad (q \in H, q' \in F). \quad (15)$$

To relax the constraint in this form, we can replace it with a penalty term in the objective that quantifies the discrepancy between demand and supply according to the

squared norm, weighed according to $\lambda > 0$:

$$\frac{1}{2\lambda} \|G^\top(q - q')\|_2^2.$$

This leads to the relaxed dual formulation in (13). We will write \mathcal{D}_λ to refer to the objective (12) with regularization term weighed by $\lambda > 0$. As λ tends to 0, the original constraint is satisfied exactly (supply matches demand), so by convention $\mathcal{D}_0 \equiv \mathcal{D}$. In the regularized primal formulation (12), $\lambda > 0$ shrinks the price parameter vector towards 0. The proof of Theorems 3.1 and 3.3 appears in Appendix A. The case of $\lambda = 0$ follows from linear programming duality. When $\lambda > 0$, general convex duality (Fenchel duality) is invoked. The duality proof admits other choices for the convex regularizer besides the squared norm, which could lead to improved bounds on convergence rates.

The following result bounds the magnitude of the optimal price parameter, which is an essential element of our convergence rate analyses.

PROPOSITION 3.4. *Consider the optimization problems in (9) and (12), and set $V = \|v\|_\infty$.*

- *If $G = G_{\text{id}}$, then every optimal \bar{w} satisfies $\|\bar{w}\|_\infty \leq (n + 1)V$ and $\|\bar{w}\|_2 \leq (n + 1)V\sqrt{\ell}$.*
- *If $G = G_{p(r)}$ for some integer $r \geq 1$, then every optimal \bar{w} satisfies $\|\bar{w}\|_\infty \leq (n + 1)V2^r$ and $\|\bar{w}\|_2 \leq (n + 1)V2^r m^{r/2}$.*

The V bound here amounts to a choice of units. It could be normalized to $V = 1$, but we choose to keep it explicit to clarify how it affects the choice of step-size in the auction. The proposition establishes that the optimal price parameter (even when non-unique) lies in a bounded set. Importantly, the size of this set depends on the number of degrees of freedom in $w \in \mathbf{R}^d$, meaning the dimension d . This implies that the price parameters w can be constrained during the auction to lie in a ball of sufficiently bounded radius without affecting the optimum. We will see that the radius of this ball affects the convergence rate of the auction; simply put, increasing the representation power of G increases the radius of the ball.

As a sketch, the proof first proceeds by bounding $\|p\|_\infty$, where recall that $p = Gw$ lists the explicit bundle prices. The reasoning is that as the components of p become large, the second term of \mathcal{D} grows, and as the components of p become small, the first term grows. Taken together, it follows that the optimum must be bounded. We then obtain bounds on $\|w\|_\infty$ given the bound on $\|p\|_\infty$. For $G = G_{\text{id}}$ this is trivial because the latter is a stack of identity matrices. The case of $G = G_{p(r)}$ requires a combinatorial argument. Finally, a bound on $\|w\|_\infty$ immediately yields a bound on $\|w\|_2$.

4. ITERATIVE AUCTION.

We now introduce the iterative auction that will be the subject of our analysis. The auction is described in Figure 1 and follows a high-level outline common to many practical auction designs: 1) prices are quoted; 2) agents place bids on bundles, meaning offers to purchase bundles at the quoted prices; 3) the seller computes a provisional allocation based on bids and prices; 4) prices are updated when there is discrepancy between the bids and allocation. Note that the agents do not make price offers, so this is a *clock auction*.

The auction is fully specified given the matrix G , the regularization parameter $\lambda \geq 0$, the step sizes η^t , and the projection radius R . Since different choices for the matrix G correspond to different pricing schemes, Figure 1 may more generally be viewed as a class of auctions. The price update in step 4 has a natural economic interpretation. Bid b^t represents demand, while allocation q^t represents supply, so $q^t - b^t$ corresponds to *excess supply* in bundle space. By applying G^\top , this maps to excess supply in feature

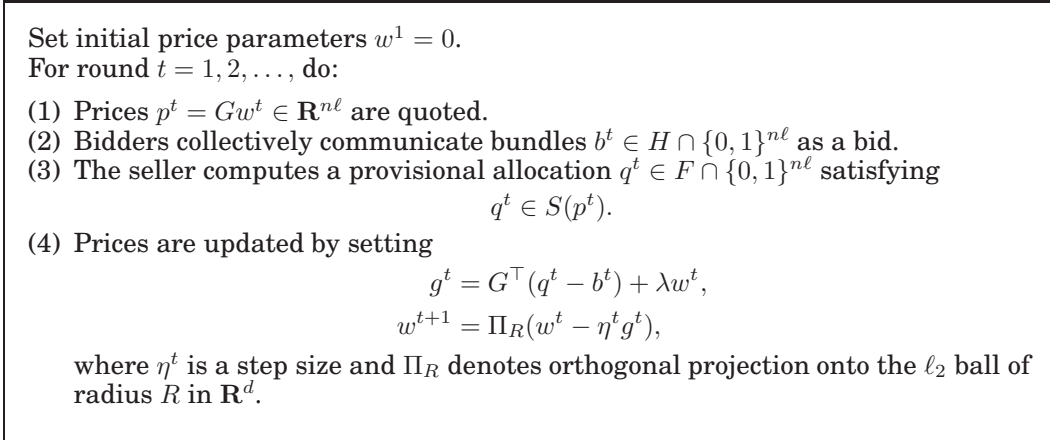


Fig. 1. Subgradient auction.

space, where the price parameter vector lies. The result is then subtracted away from the current price parameter w^t , which is therefore updated in the direction of *excess demand*. With regularization ($\lambda > 0$), the update is pulled back towards the current iterate w^t , leading to more conservative price parameter updates.

The projection in step 4 is essentially a rescaling to ensure that price coordinates are well-behaved, and we confirm below that the radius can be chosen large enough (relative to the number of agents and items) to leave the final clearing prices unaffected. Projection may seem undesirable, since small choices of R can rule out the (unconstrained) optima guaranteed to exist for the pricing problem in Theorems 3.1 and 3.3. However, as detailed in Proposition 3.4, it is possible to choose R purely from knowledge of the matrix G and units V so that the ball of radius R includes unconstrained optima. This point will rearise in Sections 5 and 6, where the iterative auction is shown to converge to these unconstrained optima under an appropriately chosen R .

Note that in step 2 of the auction, agents are required to provide an *integer* bid vector, representing a bundle with whole items. This is where item indivisibility comes into play. Under our model, requiring integer bid vectors is unrestrictive, because an agent's best-response problem is to optimize a linear function (valuation minus price) over the consumption set (1), whose extreme points are integer. Therefore, there is always an integer best-response—in plainer terms, an agent can always maximize its utility by bidding on a bundle with whole items. Similarly, the seller is required to compute an integer allocation in step 3. As prices are a linear function and the production set (2) has integer extreme points, this does not prevent the seller from computing a revenue-maximizing allocation.

It is worth noting that in our auction neither the price parameters $(w^t)_{t \geq 1}$ nor the bundle prices $(p^t)_{t \geq 1}$, where $p^t = Gw^t$, are constrained to move monotonically (ascending or descending pointwise). While most iterative auctions in the literature are monotone, there is precedent for non-monotonic price paths, for instance Ausubel's auction for multiple distinct items [Ausubel 2006]. On one hand, non-monotone price paths allow for highly adversarial bidding behavior that can even prevent the auction from converging; we return to this issue in Section 6 and show how it is addressed by activity rules. On the other hand, non-monotone auctions also have substantial advantages

when it comes to convergence rates, which makes them more natural to study in the context of this paper.⁵

We next confirm that there is a direct relationship between the optimization problems of Section 3 and the iterative auction in Figure 1. First, observe that the auction makes no mention of agent valuations—valuations affect agent bidding behavior, but cannot be part of the auction specification. We say that a bid vector $b \in H$ and valuation vector $v \in \mathbf{R}^{n\ell}$ are *consistent* with each other at prices p if $b \in U(p; v)$. In words, this means that the bid vector is a best-response to prices assuming the given valuations (i.e., the bid maximizes the agents’ aggregate utility).

PROPOSITION 4.1. *Let representation matrix G and regularization parameter $\lambda \geq 0$ be given. If bid vector $b^t \in H$ is consistent with valuation v^t at prices $p^t = Gw^t$, and allocation $q^t \in F$ maximizes the seller’s revenue at prices p^t , then*

$$g^t = G^\top(q^t - b^t) + \lambda w^t \tag{16}$$

is a subgradient of $\mathcal{D}_\lambda(\cdot; v^t)$ at w^t .

The full proof is given in the appendix. This proposition connects the optimization problem and the iterative auction; specifically, since the subgradient expression in (16) matches the update step in the auction, namely step 4 in Figure 1, the auction is performing subgradient descent on \mathcal{D}_λ . Under this viewpoint we can analyze auction convergence using well-developed techniques.

Implementation. Although our focus is on the theoretical convergence properties of the subgradient auction in Figure 1, let us say a few words about how it might be implemented in practice. This is a salient question given that the auction, as specified, manipulates high-dimensional vectors and computes allocations over a complex polytope.

Note first that although bids b^t and allocations q^t have dimension $n\ell$, and ℓ may be exponential in m , the restriction that they should be integer in steps 2 and 3 means that they have at most n non-zero entries and are succinct to communicate. The sparsity of bids and allocations also means that the price update in step 4 can be computed efficiently. As mentioned previously, the projection in the price update is simply a rescaling of the price parameters.

As the representation matrix G is commonly known, communicating prices is a matter of communicating $w^t \in \mathbf{R}^d$, which can be done directly for linear prices or polynomials of low degree. For larger d , the following result gives a dual representation of w^t with storage on the order of nt coefficients. In the statistical learning literature, such results are known as *representer theorems* [Steinwart and Christmann 2008].

PROPOSITION 4.2. *Let matrix G and regularization parameter $\lambda \geq 0$ be given. For all t , $w^t \in \text{span}(G^\top)$. Indeed, setting $\gamma^1 = 1$ and $\gamma^{t+1} = R/\max\{R, \|w^t - \eta^t g^t\|_2\}$, w^t has the form*

$$w^t = G^\top \sum_{s=1}^{t-1} \gamma^t \eta^s (b^s - q^s) \prod_{j=s+1}^{t-1} \gamma^{j+1} (1 - \lambda \eta^j),$$

meaning w^t may be reconstructed from $(\eta^s, b^s, q^s, \gamma^s)_{s=1}^{t-1}$.

⁵To illustrate, consider the task of guessing a secret number: One player chooses a number y within the interval $[0, 1]$, and another player must guess this number to within accuracy $\epsilon > 0$. The game protocol is that the guesser specifies an $x \in [0, 1]$, and the chooser responds with $\mathbf{1}[x \geq y]$. With monotonic guesses, it is necessary to make $\Omega(1/\epsilon)$ queries; on the other hand, non-monotonic questions allow binary search, meaning $\mathcal{O}(\ln(1/\epsilon))$ queries.

The proof of this statement can be found in the appendix. The claim follows by an inductive argument, since by the price update rule w^t is a linear combination of elements from $\text{span}(G^T)$.

The remaining item is to compute a revenue-maximizing allocation in step 3. The efficient allocation problem in (3) is known to be NP-hard by reduction from weighted set packing [Nisan 2000], and comparing (6) and (3), we see that the allocation step is equivalently intractable. In practice, this step is implemented using integer programming solvers, and there is a large body of work on solvers for combinatorial auctions [e.g. Cramton et al. 2006, Part III]. Using an efficient allocation oracle in the specification allows us to focus on the number of auction rounds, which is the relevant metric for iterative combinatorial auctions.

There are several parameters to tune to run the auction: the regularization weight λ , the projection radius R , and the step size schedule $\{\eta^t\}$. Our convergence rate analyses—see Theorems 5.1 and 6.1—will provide guidance on how to set each of these parameters. Essentially, all of them can be set in terms of V , the maximum possible agent value for a bundle.

The convergence results we present in the next sections concern price convergence (in the primal), not allocation convergence (in the dual), and even for prices convergence holds only in the limit rather than in a finite number of steps. For auctions based on the subgradient method, it is not possible to establish finite-time convergence without further structural assumptions [de Vries et al. 2007]. To increase the chances of matching supply with demand at each round, practical auction designs break ties in step 3 to satisfy as many agents as possible, and also allow agents to take an ϵ -discount on bundles in the provisional allocation [e.g. Parkes 1999]. For relatively small ϵ , this does not impact our bounds on convergence rate.

5. STOCHASTIC BIDDING.

In this section and the next, we provide the central results of the paper which bound the convergence rate of prices in the subgradient combinatorial auction of Figure 1. We first consider *stochastic bidding*, meaning that agents behave according to a random utility model. In the next section we turn our attention to *adversarial bidding* in which agents can place arbitrary bids, constrained across rounds only by activity rules.

Our stochastic model aims to capture the fact that bids can incorporate an element of randomness at each round due to fluctuating valuations, bounded rationality, behavioral noise, etc. However, rather than directly assume that bids are stochastic, we instead assume that valuations are stochastic at each round and that bids are chosen as best-responses to prices according to the realized valuations. This is the bidding behavior that arises from a random utility model, familiar from discrete choice modeling [McFadden et al. 1973]. Formally, at each round the agents draw their valuation profile v^t from a fixed distribution, denoted ν . That is, one should view each v^t , for $t = 1, 2, \dots$, as an i.i.d. draw from ν . Once v^t is drawn and the prices $p^t = Gw^t$ are quoted by the auction, the agents place a collective bid vector b^t consistent with v^t , where ties are broken arbitrarily in case the best-response is not unique. In discrete choice models, the random valuation is usually decomposed as $v^t = \tilde{v} + \epsilon^t$, where \tilde{v} is the mean valuation and ϵ^t is an error term capturing deviations from the mean at round t . The most common error models for ϵ^t are the Gumbel distribution (known as the logit model) and the Gaussian distribution (known as the probit model).

Under stochastic bidding, bid vector b^t is consistent with random valuation v^t under prices p^t at each round, which means that the auction is performing subgradient descent as per Proposition 4.1. Note that the associated distribution ν may be arbitrary, and in particular ν does not need to be a product distribution across the n bidders. Our convergence results are robust to arbitrary correlations between the bidders valu-

ations. However, one limitation of the model is that the i.i.d. nature of the distribution cannot incorporate learning from past bids and prices (e.g., as one would expect if there were a common value component to the agents' valuations).

THEOREM 5.1. *Let ν be a distribution over value vectors $v \in \mathbf{R}^{n\ell}$ with $V = \mathbf{E}_\nu[\|v\|_\infty] < \infty$; by these conditions, there exists an optimum \bar{w} to the problem $\min\{\mathbf{E}_\nu[\mathcal{D}_\lambda(w; v)] : w \in \mathbf{R}^d\}$. Moreover, with probability at least $1 - \delta$ over an i.i.d. draw of valuations $(v^t)_{t=1}^T$, running an iterative auction over T rounds with step size $\eta^t = V/\sqrt{t}$, regularization $\lambda \leq 1/V$, and any projection radius $R \geq \|\bar{w}\|_2$ gives the bound*

$$\begin{aligned} \mathbf{E}_\nu[\mathcal{D}_\lambda(\hat{w}^T; v)] - \mathbf{E}_\nu[\mathcal{D}_\lambda(\bar{w}; v)] &\leq \sum_{t=1}^T \hat{\eta}^t (\mathbf{E}_\nu[\mathcal{D}_\lambda(w^t; v)] - \mathbf{E}_\nu[\mathcal{D}_\lambda(\bar{w}; v)]) \\ &\leq \mathcal{O}\left(\frac{\kappa^2 \ln T \sqrt{\ln(1/\delta)} V}{\sqrt{T}}\right) \end{aligned}$$

where $\hat{\eta}^t = \eta^t / \sum_{s=1}^T \eta^s$ and $\hat{w}^T = \sum_{t=1}^T \hat{\eta}^t w^t$ is the averaged iterate. The quantity κ depends on representation matrix G and projection radius R , and may be bounded as follows.

— When $G = G_{\text{id}}$, it suffices to choose $R = (n+1)V\sqrt{\ell}$, whereby

$$\kappa \leq \sqrt{n} + \sqrt{m} + 2(n+1)\sqrt{\ell}.$$

— When $G = G_{p(r)}$, it suffices to choose $R = (n+1)Vm^{r/2}2^r$, whereby

$$\kappa \leq (1 + \sqrt{n})m^r + 2(n+1)m^{r/2}2^r.$$

There are several terms in the bound of Theorem 5.1, but the leading term κ^2 roughly reflects the number of degrees of freedom (i.e., the dimension d of the price parameter w), and we see that increasing polynomial degree or using bundle pricing weakens guarantees on convergence time. The quantity V essentially captures the scale of the bidder valuations. In the simplest case where ν has compact support, V corresponds to the largest possible value for a bundle. Lemma B.7 in the appendix provides bounds on V for the logit and probit models, as well as any error distribution with subgaussian tails; all bounds have the form $\|\bar{v}\|_\infty + \mathcal{O}(\sigma_{\max} \ln(n\ell))$, where \bar{v} is the mean valuation and σ_{\max} is the maximum over the valuation's coordinate-wise standard deviations.

It is worth stressing how the bound in Theorem 5.1 (as well as the upcoming bound in Theorem 6.1) departs from standard statistical treatments. In statistical learning theory, it is standard to choose either the radius R , or the regularization weight $\lambda > 0$, so as to provide faster or slower convergence. The same applies to the non-sequential (batch) setting, where a typical bound for kernel classifiers depends purely on $\lambda > 0$ [Boucheron et al. 2005, Corollary 4.3]. Such an approach is not possible here because R must be chosen so as to leave the set of optimal solutions intact—otherwise, the optimal solution would lose its meaning as clearing prices. The main challenge in proving Theorem 5.1 is to show how this can be achieved with bounded choices of R .

Theorem 5.1 only controls convergence of the objective function \mathcal{D}_λ , not the prices themselves. If we want to claim that the individual bundle prices during the auction are ‘informative’ to the bidders, then the coordinates of $p^t = Gw^t$ should also be stable. To this end, we provide the following result.

COROLLARY 5.2. *Consider the setting of Theorem 5.1 but with $\lambda \in (0, 1/V]$, and let $w \in \mathbf{R}^d$ be any vector satisfying $\mathbf{E}_\nu[\mathcal{D}_\lambda(w; v)] \leq \mathbf{E}_\nu[\mathcal{D}_\lambda(\bar{w}; v)] + \epsilon$ for some $\epsilon \geq 0$. Then $\|w - \bar{w}\|_2^2 \leq 2\epsilon/\lambda$. In particular, after T rounds, with probability at least $1 - \delta$ prices*

$\hat{p}^T = G\hat{w}^T$ and $\bar{p} = G\bar{w}$ satisfy

$$\|\hat{p}^T - \bar{p}\|_\infty \leq \mathcal{O}\left(\frac{\kappa\|G\|_{2,\infty}}{\lambda} \sqrt{\frac{\ln(T)\sqrt{\ln(1/\delta)}}{\sqrt{T}}}\right),$$

where κ may be bounded as in Theorem 5.1, and $\|G\|_{2,\infty} = \max\{\|Gw'\|_\infty : \|w'\|_2 \leq 1\}$ can be bounded as $\|G_{\text{id}}\|_{2,\infty} \leq 1$ and $\|G_{p(r)}\|_{2,\infty} \leq m^{r/2}$.

In words, this statement converts the convergence in objective value from Theorem 5.1 to convergence in prices themselves, assuming $\lambda > 0$. However, Corollary 5.2 has a few weaknesses: (1) $\lambda > 0$ must be chosen small to ensure there is not too much discrepancy between demand and supply, thus the bound converges slowly; (2) the economic meaning of regularizing \bar{w} is still unclear—it appears to favor a bidder-optimal choice of prices, but we have no formal statements to this effect. Note that the choice $\lambda = 1/V$ causes the right hand side to scale linearly with V , matching Theorem 5.1 and also the interpretation of V as units or scale.

Bias-Variance. The results of this section can be interpreted in terms of a bias-variance trade-off. The infimal value of $\mathbf{E}_\nu[\mathcal{D}_\lambda(\cdot; v)]$ depends on the representation matrix G , and in this way represents the ‘bias’ of the auction. Using a more expressive class of prices reduces this bias, and the lowest bias is attained when the prices can support an integer optimal solution. On the other hand, the actual value of the bounds, which gives the rate of convergence of the auction, is the ‘variance’ term. Increasing price expressiveness weakens the bound, and therefore simpler pricing matrices (e.g., polynomial matrices of low degree) should exhibit faster convergence. One way to moderate the trade-off is to use the simplest class of prices available (in terms of dimensionality) that clears the market, although this can be hard to know a priori.

6. ARBITRARY BIDDING.

In this section we turn to a model where agent bids can be essentially arbitrary, and even adversarial across rounds. While the model is behaviorally unreasonable without further constraints, it is still possible to provide a certain convergence guarantee on the objective value. On the other hand, it is not possible to obtain price convergence: the optimal price vectors may drift and oscillate, in contrast with the conclusions we were able to draw in Corollary 5.2. The model therefore motivates the use of *activity rules* to constrain agent bids across rounds, and this section shows how a well-designed activity rule can result in meaningful convergence guarantees for both the objective and prices.

The arbitrary bidding model is as follows. In round t , upon seeing prices $p^t = Gw^t$, bidders collectively release an integer bid vector $b^t \in H \cap \{0, 1\}^{n^\ell}$ as specified in step 2 of the auction. In contrast with the stochastic model of the previous section, where b^t must be a best-response with respect to the random valuation vector v^t according to (5), b^t need only be consistent with some valuation vector v^t . Now, without any constraints on the space of valuations, one can always find a valuation vector with which a given bid vector is consistent, whatever the prices. The first result of this section merely requires there to exist choices of $(v^t)_{t \geq 1}$ which satisfy $\|v^t\|_\infty \leq V$ for some scalar V . This is a mild constraint which effectively means that $b_i^t(x)$ must be zero whenever $p_i^t(x)$ exceeds V .

Our initial result under arbitrary bidding is the following. Superficially, the statement appears similar to the convergence statement for the stochastic model given in Theorem 5.1. The essential difference is that the left-hand side is no longer competing with a fixed target $\inf_w \mathbf{E}_\nu[\mathcal{D}_\lambda(w; v)]$. Instead, the comparison is against

$\inf_w \sum_{t \in [T]} \mathcal{D}_\lambda(w; v^t)$, and each term in the summation can change drastically at each round t . This objective does have an economic interpretation if one views our procedure as a sequential posted price mechanism rather than an iterative auction. A new set of n bidders arrives at each round, and a new set of m items is available for sale. The seller's problem is to try to post prices that clear the market at each round t , before bidder valuations v^t are revealed, where clearing quality is captured by the objective $\mathcal{D}_\lambda(\cdot; v^t)$. The result bounds the *regret* of the procedure against the best fixed prices in hindsight, which is a standard objective for online algorithms [Hazan 2012].

THEOREM 6.1. *Consider an iterative auction where bid vectors $(b^t)_{t=1}^T$, with some consistent sequence of value vectors $(v^t)_{t=1}^T$, are announced in alternation with price parameters $(w^t)_{t=1}^T$ provided by the auction mechanism invoked with step size $\eta^t = V/\sqrt{t}$ for some $V \geq 0$, regularization $\lambda \leq 1/V$, and some projection radius $R \geq 0$. Then there exists a minimizer \bar{w}^T to $f(w) = \sum_{t=1}^T \eta^t \mathcal{D}_\lambda(w; v^t)$, and if $V \geq \sup_{t \in [T]} \|v^t\|_\infty$ and $R \geq \|\bar{w}^T\|_2$, then*

$$\sum_{t=1}^T \hat{\eta}^t (\mathcal{D}_\lambda(w^t; v^t) - \mathcal{D}_\lambda(\bar{w}^T; v^t)) \leq \mathcal{O}\left(\frac{\kappa^2 V \ln T}{\sqrt{T}}\right),$$

where $\hat{\eta}^t = \eta^t / \sum_{s=1}^T \eta^s$. The κ quantity depends on representation matrix G and projection radius R , and may be bounded as in Theorem 5.1.

The proof is very similar to that of Theorem 5.1. Comparing this bound to Theorem 5.1, nearly everything is the same, including the general growth of the leading term κ^2 in response to choosing G_{id} or $G_{p(r)}$. However, as mentioned above, what differs is the left-hand term: progress is measured against a time-varying target rather than a time-independent target as in the stochastic model. As the optimal pricing vector \bar{w}^t is a function of time, it need not converge in any way, and we cannot hope for convergence in prices either. To illustrate this concretely, we have the following result.

PROPOSITION 6.2. *Suppose the setting of Theorem 6.1 with $n = 2$ bidders and $m = 1$ item, but with step sizes $(\eta^t)_{t \geq 1}$ being any positive reals satisfying $\sum_{t \geq 1} \eta^t = \infty$. Under bundle or polynomial prices (of any degree), there exists a bidding sequence $(b^t)_{t \geq 1}$ consistent with a valuation sequence $(v^t)_{t \geq 1}$ such that every sequence of optimal price parameters $(\bar{w}^t)_{t \geq 1}$ fails to converge.*

The proof of this fact, given in Appendix A, constructs a concrete bidding sequence whereby the corresponding optima $(\bar{w}^t)_{t \geq 1}$ oscillate between two cluster points.

To link the behavior of bidders across rounds and recover price convergence, the auction can make use of *activity rules*. In fact, activity rules are used in practice specifically to disallow certain kinds of adversarial bidding behaviors that are assumed away by simple models of best-response agents (e.g., bid parking and sniping) [Ausubel and Baranov 2014a]. The rule that is most firmly grounded in theory is the *revealed preference activity rule*, also called GARP (after the *generalized axiom of revealed preference*) [Ausubel and Baranov 2014b]. We consider here the strictest form of the rule which requires exact adherence to the GARP axiom. A sequence of bid vectors $(b^t)_{t \geq 1}$, placed in response to prices $(p^t)_{t \geq 1}$ over rounds, satisfies the *GARP activity rule* if for every sequence of distinct rounds $t_1, t_2, \dots, t_{k'}$ (not necessarily consecutive or ordered),

$$\sum_{k=1}^{k'} (b^{t_{k+1}} - b^{t_k})^\top p^{t_k} \geq 0, \quad (17)$$

with the convention $t_{k'+1} = t_1$. Our analysis can also accommodate more relaxed forms where bidders are allowed to violate GARP in earlier rounds. The GARP activity rule can be enforced efficiently in practice using network flow algorithms [Vohra 2004].

A sequence of bid vectors satisfies the GARP activity rule if and only if there is a fixed valuation vector v consistent with the entire sequence of bids and prices (even though the bidders may not be explicitly considering such a valuation); for completeness, a proof of this fact is provided as Lemma B.6 in the appendix. Under these circumstances, Theorem 6.1 can be strengthened to obtain a more meaningful bound.

THEOREM 6.3. *Consider the setting of Theorem 6.1, and assume the bid sequence $(b^t)_{t \geq 1}$ satisfies the GARP activity rule with respect to prices $(p^t)_{t \geq 1}$. Then there exists a single value vector v that is consistent with b^t under prices p^t , for all rounds t , and moreover*

$$\mathcal{D}_\lambda(\hat{w}^T; v) - \mathcal{D}_\lambda(\bar{w}; v) \leq \mathcal{O}\left(\frac{\kappa^2 V \ln T}{\sqrt{T}}\right),$$

where $\hat{w}^T = \sum_{t=1}^T \eta^t w^t / \sum_{s=1}^T \eta^s$ is the averaged iterate, \bar{w} is a minimum for $\mathcal{D}_\lambda(\cdot; v)$, and κ may be bounded as in Theorem 6.1. Moreover, if $\lambda \in (0, 1/V]$, then

$$\|\hat{p}^T - \bar{p}\|_\infty \leq \mathcal{O}\left(\frac{\kappa \|G\|_{2,\infty}}{\lambda} \sqrt{\frac{\ln T}{\sqrt{T}}}\right).$$

where $\hat{p}^T = G\hat{w}^T$ and $\bar{p} = G\bar{w}$, and $\|G\|_{2,\infty}$ may be bounded as in Corollary 5.2.

In simpler terms, the GARP activity rule ensures that bidding across rounds is consistent with at least one fixed valuation profile, and for any such profile the auction converges in objective value. With regularization ($\lambda > 0$), we also obtain price convergence, in contrast to Proposition 6.2.

7. CONCLUSION.

This paper obtained concrete bounds on the rate of convergence of iterative auctions that correspond to subgradient methods for the underlying optimization problem over prices. Our setup can accommodate many different pricing schemes and allows one to analyze and compare them under a single framework. It also admits bidder behaviors beyond straightforward best-response bidding. We considered two generalizations of straightforward bidding: stochastic bidders and arbitrary bidders, with the latter constrained by activity rules.

The convergence rates obtained under both models are very similar. In both cases, using a more expressive pricing scheme weakens convergence guarantees. Bounds are proportional to the degrees of freedom in the prices, so item price bounds are exponentially better than bundle price bounds. Our analysis quantifies one side of the trade-off between convergence rate and ability to clear the market, which mirrors the bias-variance trade-off familiar from statistics. Our results suggest that an iterative combinatorial auction should use the simplest class of prices possible that can clear the market. In this respect, polynomial prices may prove useful in practice.

The stochastic model of bidders allows for errors in their estimates of the values for various bundles in each round. This translates into stochastic and imperfect best-response behavior, following standard random utility models. Our convergence results show that subgradient auctions perform effective price discovery even under bidding errors, and the results are robust to correlations between the bidders' valuations. With regularization, we also obtain convergence in individual bundle prices. Under the arbitrary bidding model, prices can oscillate and fail to converge. This motivates the use of revealed preference activity rules (GARP), which restore consistency with a fixed

valuation vector. Our analysis draws a connection between the constraints imposed by the activity rule and convergence of the auction.

There are many ways to extend and build on our analysis. An important avenue for future work is to derive lower bounds on convergence rates and thereby achieve a separation between pricing schemes. There are standard tools for producing lower bounds on subgradient methods [e.g. Agarwal et al. 2009], but they might involve pathological constructions lacking economic meaning. Another avenue is to study other pricing schemes besides polynomial or bundle pricing, which could be relevant in specific domains and worthy of study. Given their prominence in practice, it would also be worthwhile to analyze monotone price auctions (ascending or descending) within the present framework. Monotone price auctions can be obtained by modifying the projection operation in the subgradient method.

ACKNOWLEDGMENTS

The authors would like to thank Jason Hartline for valuable feedback. We also thank participants at seminars at Duke, the University of Zurich, and the 2015 AMMA Conference for comments.

REFERENCES

- Alekh Agarwal, Martin Wainwright, Peter Bartlett, and Pradeep Ravikumar. 2009. Information-theoretic lower bounds on the oracle complexity of convex optimization. In *Advances in Neural Information Processing Systems*. 1–9.
- Kenneth Arrow and Frank Hahn. 1971. *General Competitive Analysis*. Holden-Day, San Francisco, CA. Now distributed by North-Holland, Amsterdam.
- Lawrence M. Ausubel. 2006. An efficient dynamic auction for heterogeneous commodities. *The American Economic Review* (2006), 602–629.
- Lawrence M. Ausubel and Oleg V. Baranov. 2014a. Market Design and the Evolution of the Combinatorial Clock Auction. *The American Economic Review* 104, 5 (2014), 446–451.
- Lawrence M. Ausubel and Oleg V. Baranov. 2014b. *A Practical Guide to the Combinatorial Clock Auction*. Technical Report. University of Maryland.
- Lawrence M. Ausubel and Peter Cramton. 2011. *Activity rules for the combinatorial clock auction*. Technical Report. University of Maryland.
- Lawrence M. Ausubel and Paul R. Milgrom. 2002. Ascending auctions with package bidding. *Advances in Theoretical Economics* 1, 1 (2002).
- Amir Beck and Marc Teboulle. 2003. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters* 31, 3 (2003), 167–175.
- Edward A. Bender and Jay R. Goldman. 1975. On the applications of Möbius inversion in combinatorial analysis. *Amer. Math. Monthly* (1975), 789–803.
- Dimitri P. Bertsekas. 1998. *Network Optimization: Continuous and Discrete Models*. Athena Scientific.
- Sushil Bikhchandani, Sven de Vries, James Schummer, and Rakesh V. Vohra. 2002. Linear programming and Vickrey auctions. *Mathematics of the Internet: E-auction and Markets* 127 (2002), 75–116.
- Sushil Bikhchandani and Joseph M. Ostroy. 2002. The package assignment model. *Journal of Economic Theory* 107, 2 (2002), 377–406.
- Jonathan M. Borwein and Adrian S. Lewis. 2010. *Convex analysis and nonlinear optimization*. Springer Science & Business Media.
- Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. 2005. Theory of Classification: A Survey of Recent Advances. *ESAIM: Probability and Statistics* 9 (2005), 323–375.

- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. 2013. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford.
- Sébastien Bubeck. 2014. Theory of Convex Optimization for Machine Learning. (2014). arXiv:1405.4980 [math.OC].
- Nicolo Cesa-Bianchi and Gábor Lugosi. 2006. *Prediction, learning, and games*. Cambridge University Press.
- Peter Cramton, Yoav Shoham, and Richard Steinberg. 2006. *Combinatorial auctions*. MIT press.
- Vincent P. Crawford and Elsie Marie Knoer. 1981. Job matching with heterogeneous firms and workers. *Econometrica* (1981), 437–450.
- Robert Day and Peter Cramton. 2012. Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research* 60, 3 (2012), 588–603.
- Sven de Vries, James Schummer, and Rakesh V. Vohra. 2007. On ascending Vickrey auctions for heterogeneous objects. *Journal of Economic Theory* 132, 1 (2007).
- Ana Fostel, Herbert E. Scarf, and Michael J. Todd. 2004. Two new proofs of Afriat’s theorem. *Economic Theory* 24, 1 (2004), 211–219.
- Pavithra Harsha, Cynthia Barnhart, David C. Parkes, and Haoqi Zhang. 2010. Strong activity rules for iterative combinatorial auctions. *Computers & Operations Research* 37, 7 (2010), 1271–1284.
- Elad Hazan. 2012. The Convex Optimization Approach to Regret Minimization. *Optimization for machine learning* (2012), 287.
- Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. 2001. *Fundamentals of Convex Analysis*. Springer Publishing Company.
- Alexander S. Kelso and Vincent P. Crawford. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica* (1982), 1483–1504.
- Sébastien Lahaie. 2009. A Kernel method for Market Clearing. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*. 208–213.
- Sébastien Lahaie. 2011. A Kernel-Based Iterative Combinatorial Auction. In *Proceedings of the 25th National Conference on Artificial Intelligence (AAAI)*. 695–700.
- Daniel McFadden and others. 1973. *Conditional logit analysis of qualitative choice behavior*. Institute of Urban and Regional Development, University of California.
- Paul R. Milgrom. 2004. *Putting auction theory to work*. Cambridge University Press.
- Noam Nisan. 2000. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*. ACM, 1–12.
- David C. Parkes. 1999. *iBundle: An efficient ascending price bundle auction*. In *Proceedings of the 1st ACM Conference on Electronic Commerce*. ACM, 148–157.
- R. Tyrrell Rockafellar. 1970. *Convex Analysis*. Princeton University Press.
- Tobias Scheffel, Alexander Pikovsky, Martin Bichler, and Kemal Guler. 2011. An experimental comparison of linear and nonlinear price combinatorial auctions. *Information Systems Research* 22, 2 (2011), 346–368.
- Stefan Schneider, Pasha Shabalin, and Martin Bichler. 2010. On the robustness of nonlinear personalized price combinatorial auctions. *European Journal of Operational Research* 206, 1 (2010), 248–259.
- Shai Shalev-Shwartz. 2011. Online learning and online convex optimization. *Foundations and Trends in Machine Learning* 4, 2 (2011), 107–194.
- Ingo Steinwart and Andreas Christmann. 2008. *Support Vector Machines*. Springer.
- Hal R. Varian. 1981. Dynamical systems with applications to economics. *Handbook of mathematical economics* 1 (1981), 93–110.
- Rakesh V. Vohra. 2004. *Advanced mathematical economics*. Routledge.
- Martin Zinkevich. 2003. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning*. 928–936.

APPENDIX

A. DEFERRED PROOFS.

This appendix provides complete proofs for all results in the paper. Several results make use of stand-alone auxiliary lemmas given in the next section.

PROOF OF THEOREMS 3.1 AND 3.3. The general existence of an optimum \bar{w} (for any $\lambda \geq 0$) follows from Lemma B.3, taking ν as a dirac measure supported on v and nowhere else. When $\lambda > 0$, the uniqueness of \bar{w} holds because the objective is strictly convex. Note that F is a polyhedral convex set. The second duality then follows by applying Lemma B.2 with $r(x) = \lambda \|x\|_2^2/2$, which gives $r^*(y) = \|y\|_2^2/(2\lambda)$ by direct computation. The first duality also follows from Lemma B.2, but now with $r(x) = 0$, meaning $r^*(y) = \iota_{\{0\}}(y)$ and $r^*(G^\top q') = \iota_{\ker(G^\top)}(q')$, which gives the dual problem as

$$\max \{v^\top q' : q' \in H, q \in F, q' - q \in \ker(G^\top)\},$$

the result following by collapsing q into the constraint on q' . The invocation of Lemma B.2 also grants the desired optimality conditions; in the case of Theorem 3.3, this directly gives the result after translating r and r^* as above, whereas with Theorem 3.1, the provided optimality conditions now give

$$\bar{q}' - \bar{q} \in \ker(G^\top), \quad \bar{q}' \in U(G\bar{w}; v), \quad \bar{q} \in S(G\bar{w}),$$

where the first can be written $\bar{q}' \in \{\bar{q}\} + \ker(G^\top)$, thus once again \bar{q} and the condition $\bar{q} \in S(G\bar{w})$ can be collapsed in by writing $\bar{q}' \in S(G\bar{w}) + \ker(G^\top)$ as desired. \square

PROOF OF COROLLARY 3.2. If G is full rank, then $\ker(G^\top) = \{0\}$, and thus the constraint in the dual objective (10) collapses to

$$q \in H \cap (\ker(G^\top) + F) = H \cap F = F.$$

Recalling (2), F is a polytope with integer extreme points corresponding to allocations of whole items. Thus the dual linear program has an integer optimal solution. The optimality conditions (11) collapse to $\bar{q} \in U(\bar{p}) \cap S(\bar{p})$ where $\bar{p} = G\bar{w}$, which shows that the dual optimal solution leads to supporting prices. \square

PROOF OF PROPOSITION 3.4. Let v be given as specified, and let ν be a dirac measure supported on v and nowhere else. The general existence of an optimum \bar{w} (for any $\lambda \geq 0$) follows from Lemma B.3. When $\lambda > 0$, uniqueness of \bar{w} is a consequence of strict convexity of the objective. The bounds on \bar{w} under various choices of G are provided by Lemma B.4, choosing ν as before. \square

PROOF OF PROPOSITION 4.1. The first statement is a consequence of standard subgradient rules [Hiriart-Urruty and Lemaréchal 2001, Theorem D.2.2.1, Theorem D.4.1.1, Theorem D.4.2.2], and the definition of subgradient descent is also standard [Bubeck 2014]. The second statement is from Lemma B.4. \square

PROOF OF PROPOSITION 4.2. The expression for w^t holds in the case $w^1 = 0$, and by induction and the definition of w^{t+1} in the mechanism,

$$\begin{aligned} w^{t+1} &= \gamma^{t+1} (w^t - \eta^t g^t) \\ &= \gamma^{t+1} (1 - \lambda \eta^t) \left(G^\top \sum_{s=1}^{t-1} r^s \eta^s (b^s - q^s) \prod_{j=s+1}^{t-1} \gamma^{j+1} (1 - \lambda \eta^j) \right) + \gamma^{t+1} \eta^t G^\top (b^t - q^t) \\ &= G^\top \sum_{s=1}^t \gamma^{t+1} \eta^s (b^s - q^s) \prod_{j=s+1}^t \gamma^{j+1} (1 - \lambda \eta^j) \end{aligned}$$

as desired. \square

PROOF OF THEOREM 5.1. First note that an optimal \bar{w} exists by Lemma B.3. Let g^t be defined as in (16), and let δf^t be any subgradient of $\mathcal{D}_\lambda(\cdot; v^t)$ at w^t . Define the quantities

$$L = \max_{t \in [T]} \left\{ \max\{\|g^t\|_2, \|\delta f^t\|_2\} \right\} \quad \text{and} \quad B = \max_{t \in [T]} \left\{ \|\bar{w}\|_2, \max_{t \in [T]} \|w^t\|_2 \right\}.$$

For convenience, define $f(w) = \mathbf{E}_\nu[\mathcal{D}_\lambda(w, v)]$. For every $t \in [T]$, by Cauchy-Schwarz and the triangle inequality,

$$(\delta f^t - g^t)^\top (w^t - \bar{w}) \leq \|\delta f^t - g^t\|_2 \|w^t - \bar{w}\|_2 \leq 4LB.$$

Consequently, since $\mathbf{E}_\nu[\delta f^t - g^t] = 0$, Azuma-Hoeffding grants, with probability at least $1 - \delta$,

$$\sum_{t=1}^T \hat{\eta}^t (\delta f^t - g^t)^\top (w^t - \bar{w}) \leq 4LB \sqrt{2 \ln(1/\delta) \sum_{t=1}^T (\hat{\eta}^t)^2} \leq 4LBV \sqrt{2(1 + \ln(t)) \ln(1/\delta)}.$$

Plugging this into Lemma B.5 and simplifying the left hand side via convexity gives the result.

For the bounds on B and L , first note that \bar{w} exists by Lemma B.3. Now consider the case $G = G_{\text{id}}$. By Proposition 4.1,

$$\|g^t\|_2 \leq \|G^\top (b^t - q^t)\|_2 + \lambda \|w^t\|_2 \leq \|b^t\|_2 + \|q^t\|_2 + \lambda B \leq \sqrt{n} + \sqrt{m} + \lambda B.$$

Similarly, taking $b \in U(p^t; v)$ and $q \in S(p^t)$, where $p^t = Gw^t$, to denote subgradient terms for v drawn from ν [Hiriart-Urruty and Lemaréchal 2001; Rockafellar 1970],

$$\|\delta f^t\|_2 \leq \|G^\top \mathbf{E}_\nu(b - q)\|_2 + \lambda \|w^t\|_2 \leq \sqrt{n} + \sqrt{m} + \lambda B.$$

Lastly, the form of B and bound $\|\bar{w}\|_2 \leq B$ are provided by Lemma B.4.

Now consider the case $G = G_{p(r)}$. By Proposition 4.1,

$$\|g^t\|_2 \leq \|G^\top (b^t - q^t)\|_2 + \lambda \|w^t\|_2 \leq \|G^\top b^t\|_2 + \|G^\top q^t\|_2 + \lambda B \leq m^r \sqrt{n} + m^r + \lambda B.$$

The derivation of the remaining properties is as for G_{id} , and the result follows. \square

PROOF OF COROLLARY 5.2. For convenience, define $f(w) = \mathbf{E}_\nu[\mathcal{D}_\lambda(w, v)]$, and let δf be a subgradient of f at w . By first order optimality conditions and strong convexity,

$$\epsilon \geq f(w) - f(\bar{w}) \geq (w - \bar{w})^\top \delta f + \frac{\lambda}{2} \|w - \bar{w}\|_2^2 \geq \frac{\lambda}{2} \|w - \bar{w}\|_2^2,$$

which gives the first result after rearrangement. Next, the definition of $\|G\|_{2,\infty}$ gives

$$\|\hat{p}^T - \bar{p}\|_\infty = \|G\hat{w}^T - G\bar{w}\|_\infty \leq \|G\|_{2,\infty} \|\hat{w}^T - \bar{w}\|_2,$$

which gives the second bound after combining with the preceding bound and invoking Theorem 5.1 to provide the value of ϵ , and using the condition $\lambda \leq 1/V$, whereby $V \leq 1/\lambda$. Lastly, to bound $\|G\|_{2,\infty}$ for G_{id} and $G_{p(r)}$, first note that the Cauchy-Schwarz inequality implies

$$\|G\|_{2,\infty} = \max\{\|G_i(x)\|_2 : i \in [n], x \in X\};$$

consequently, in either case, we only need to count the number of 1s in rows of G . For G_{id} , this immediately gives $\|G\|_{2,\infty} \leq 1$. For $G_{p(r)}$, the bundle \tilde{x} which contains all m items will correspond to a row of 1s; since all rows have 0s and 1s, this bundle attains the maximum norm (and does not vary with bidder, so we may consider the first bidder, thus row $G_1(\tilde{x})$), which gives

$$\|G_{p(r)}\|_{2,\infty} = \|G_1(\tilde{x})\|_2 = \sqrt{\sum_{k=1}^r \binom{m}{k}} \leq \sqrt{m^r}.$$

□

PROOF OF THEOREM 6.1. An optimal \bar{w} exists by Lemma B.3, with measure ν chosen to be the discrete measure over $(v_t)_{t=1}^T$. The rate follows from Lemma B.5 with $f^t = \mathcal{D}_\lambda(\cdot, v^t)$ (whereby $\delta f^t - g^t = 0$). The estimates on L and B are as in the proof of Theorem 5.1. □

PROOF OF PROPOSITION 6.2. First note that for $m = 1$, the price parameter vector is one-dimensional (recalling our convention that the empty bundle has price of 0) and coincides for bundle and polynomial prices of all degrees. Thus we write $p = w$ for the price of the single item.

This proof will construct a sequence of bids, organized into epochs ending at times t_1, t_2, \dots , such that for any $k \geq 1$ we have $\|\bar{w}^{t_k} - \bar{w}^{t_{k+1}}\|_2 = |\bar{w}^{t_k} - \bar{w}^{t_{k+1}}| \geq 1$; since this happens for arbitrarily large choices of k , the sequence of optima is not a Cauchy sequence (and thus does not converge). The bidding behavior will be defined in terms of valuations v^t at time t , and the bids b^t (which will not require more discussion in this construction) are merely any choice which maintains the consistency of the valuations.

The construction is as follows. In every round, both bidders have the same values, and moreover assign value 0 to the empty bundle. In even epochs, both assign value 1 to the item, whereas in odd epochs they assign it value 0.

The epoch lengths $(t_k)_{k \geq 0}$ will be constructed so that the objective functions places more emphasis one selecting the single item in even epochs, whereas odd epochs will emphasize selecting no item. To this end, collect all the step sizes η^s from even and odd epochs into η_e^s and η_o^s , meaning

$$\eta_e^s = \sum_{i=1}^s \eta^i \mathbb{1}[i \text{ in even epoch}], \quad \eta_o^s = \sum_{i=1}^s \eta^i \mathbb{1}[i \text{ in odd epoch}].$$

Now define $(t_k)_{k \geq 0}$ inductively as $t_0 = 0$, and thereafter, given $(t_j)_{j=1}^k$, define t_{k+1} to be an integer sufficiently large so that $\eta_e^{t_{k+1}} > \eta_o^{t_{k+1}}$ when $k+1$ is even and otherwise $\eta_o^{t_{k+1}} < \eta_e^{t_{k+1}}$, where the existence of t_{k+1} is guaranteed from positivity of η^s and $\sum_{s \geq 1} \eta^s = \infty$.

Now fix some epoch k , and set $\eta_e = \eta_e^{t_k}$ and $\eta_o = \eta_o^{t_k}$ for convenience. The objective function evaluates to

$$\begin{aligned} \sum_{s=1}^{t_k} \eta^s \mathcal{D}_\lambda(w; v^s) &= 2\eta_e \max\{1-p, 0\} + 2\eta_o \max\{-p, 0\} + (\eta_e + \eta_o) \max\{p, 0\} \\ &= \begin{cases} 2\eta_e - 2(\eta_e + \eta_o)p & \text{if } p \leq 0, \\ (\eta_e + \eta_o)p & \text{if } p \geq 1, \\ 2\eta_e + (\eta_o - \eta_e)p & \text{if } p \in [0, 1]. \end{cases} \end{aligned}$$

As $\eta_e + \eta_o > 0$, we see that in the range $p \leq 0$ the minimum is uniquely reached at 0, and in the range $p \geq 1$ the minimum is uniquely reached at 1. Thus we restrict our attention to the range $p \in [0, 1]$. There, we see that if $\eta_o > \eta_e$, the unique minimum is 0, while if $\eta_o < \eta_e$, the unique minimum is 1. Thus, at the end of epoch k , we have $\bar{w}^{t_k} = 0$ if k is odd and $\bar{w}^{t_k} = 1$ if k is even. This completes the proof. \square

PROOF OF THEOREM 6.3. The existence of v is granted by Lemma B.6. Plugging in the consistent sequence $(v^t)_{t \geq 1}$ with $v^t = v$ into Theorem 6.1 gives the first inequality after collapsing the left hand side via Jensen's inequality as in the proof of Theorem 5.1. The second inequality is proved analogously to Corollary 5.2. \square

B. TECHNICAL LEMMAS.

This appendix provides results relating to the optimization problem of pricing studied throughout. The first result is a helper lemma bounding indirect utilities. The next is a generic duality result. We prove duality using a general convex regularizer, which covers the squared norm regularizer used in the paper. The two subsequent results study the minimizers of the pricing problem. We also provide a standard convergence bound for online subgradient descent, a proof that the GARP activity rule implies that bidding is consistent with a fixed valuation vector, and lastly bounds on the quantity $\mathbf{E}_\nu(\|v\|_\infty)$ of relevance to the analysis of stochastic bidding.

LEMMA B.1. *Both u and s are convex, closed, polyhedral, and nonnegative. Moreover, for any prices $p \in \mathbf{R}^{n^\ell}$ and any valuations $v \in \mathbf{R}^{n^\ell}$,*

$$u(p; v) \geq -\|v\|_\infty - \min_{\substack{i \in [n] \\ x \in X}} p_i(x) \quad \text{and} \quad s(p) \geq \max_{\substack{i \in [n] \\ x \in X}} p_i(x).$$

PROOF. Both u and s are convex, closed, and polyhedral by definition [Rockafellar 1970, Chapter 19]. They are nonnegative by their definition, noting that $0 \in F$ and $0 \in H$. The remaining inequalities follow similarly by definition: letting $q_1 \in H$ and $q_2 \in F$ denote standard basis vectors so that

$$p^\top q_1 = \min_{\substack{i \in [n] \\ x \in X}} p_i(x) \quad \text{and} \quad p^\top q_2 = \max_{\substack{i \in [n] \\ x \in X}} p_i(x),$$

then

$$u(p; v) = \max_{q \in H} (v - p)^\top q \geq v^\top q_1 - p^\top q_1 \geq -\|v\|_\infty - p^\top q_1 \quad \text{and} \quad s(p) = \max_{q \in F} p^\top q \geq p^\top q_2.$$

\square

LEMMA B.2. *Let closed convex bounded below $r : \mathbf{R}^d \rightarrow \mathbf{R}$, matrix $G \in \mathbf{R}^{n^\ell \times d}$ and vector $v \in \mathbf{R}^{n^\ell}$ be given. Then*

$$\inf \{u(Gw; v) + s(Gw) + r(w) : z \in \mathbf{R}^d\} = \max \{v^\top q' - r^*(G^\top(q' - q)) : q' \in H, q \in F\},$$

where $r^*(y) = \sup_w (y^\top w - r(w))$ is the convex (Fenchel) conjugate of r [Rockafellar 1970, Chapter 12]. Feasible points \bar{w} and (\bar{q}, \bar{q}') are optimal for the primal and dual problems, respectively, if they satisfy the conditions $G^\top(\bar{q}' - \bar{q}) \in \partial r(\bar{w})$, $\bar{q}' \in U(G\bar{w}; v)$, and $\bar{q} \in S(G\bar{w})$.

PROOF. For convenience, let ι_S denote the indicator function for convex set S , defined as $\iota_S(q) = 0$ for $q \in S$ and $+\infty$ otherwise. Additionally, we write $u(Gw) = u(Gw; v)$ since $v \in \mathbf{R}^{n_\ell}$ is fixed throughout. By Lemma B.1, both u and s are convex, closed, polyhedral, and bounded below by 0. Since $u(p) = \sup_{q \in H} (v - p)^\top q$, applying standard conjugacy rules [Rockafellar 1970, Theorem 12.3],

$$u^*(q) = \iota_H(-q) + v^\top q.$$

Similarly, as $s(p) = \sup_{q \in F} p^\top q$, we have

$$s^*(q) = \iota_F(q).$$

Combining these pieces, both u^* and s^* are polyhedral [Rockafellar 1970, Theorem 19.2, Corollary 19.2.1, Theorem 19.4], and thus

$$(u + s)^*(q) = \min \{u^*(q') + s^*(q - q') : q' \in \mathbf{R}^{n_\ell}\} \quad (18)$$

where attainment on the right-hand side holds because the conjugate is proper, which in turn holds because u and s are bounded below and F is nonempty [Rockafellar 1970, Theorem 16.4, Corollary 19.3.4].

Since u , r and s are bounded below and finite everywhere, it follows by Fenchel duality [Borwein and Lewis 2010, Theorem 3.3.5, Exercise 3.3.9.f] that

$$\inf \{u(Gw) + s(Gw) + r(w) : w \in \mathbf{R}^d\} = \max \{-(u + s)^*(-q) - r^*(G^\top q) : q \in \mathbf{R}^{n_\ell}\},$$

and moreover that a pair (\bar{w}, \bar{q}) is optimal for the primal and dual problems, respectively, iff $G^\top \bar{q} \in \partial r(\bar{w})$ and $-\bar{q} \in \partial(u + s)(G\bar{w})$.

To simplify the dual expression, note by (18) and other conjugacy relations above that

$$\begin{aligned} & \max \{-(u + s)^*(-q) - r^*(G^\top q) : q \in \mathbf{R}^{n_\ell}\} \\ &= \max \{-\min \{u^*(q') + s^*(-q - q') : q' \in \mathbf{R}^{n_\ell}\} - r^*(G^\top q) : q \in \mathbf{R}^{n_\ell}\} \\ &= \max \{-\iota_H(-q') - v^\top q' + \iota_F(-q - q') - r^*(G^\top q) : q, q' \in \mathbf{R}^{n_\ell}\} \\ &= \max \{-v^\top q' - r^*(G^\top q) : q \in \mathbf{R}^{n_\ell}, q' \in -H, q + q' \in -F\}. \end{aligned}$$

Combining this with the cosmetic changes of variable $q' \mapsto -q'$ and subsequently $q \mapsto q' - q$ leads to

$$\max \{v^\top q' - r^*(G^\top (q' - q)) : q' \in H, q \in F\}$$

as desired.

It remains to prove the sufficient conditions for optimality. Consequently, suppose $(\bar{w}, (\bar{q}, \bar{q}'))$ are given as in the statement, meaning $G^\top(\bar{q}' - \bar{q}) \in \partial r(\bar{w})$, $\bar{q}' \in U(G\bar{w})$, and $\bar{q} \in S(G\bar{w})$. In order to show these are optimal, they will be shown to satisfy both the optimality conditions above, which provide optimality of (\bar{w}, \bar{q}) , and also an optimality condition on the infimal convolution, which in turn grants optimality of \bar{q}' . To proceed with this analysis, it is necessary to first reverse the change of variable on these assumed conditions, meaning first performing $\bar{q}' - \bar{q} \mapsto \bar{q}$ and then $-\bar{q}' \mapsto \bar{q}'$, which means the transformed variables satisfy the conditions $G^\top \bar{q} \in \partial r(\bar{w})$, $\bar{q}' \in -U(G\bar{w})$, and $\bar{q} + \bar{q}' \in -S(G\bar{w})$.

According to the duality statements above, in order for (\bar{w}, \bar{q}) to be optimal, it suffices (as above) to show $G^\top \bar{q} \in \partial r(\bar{w})$ and $-\bar{q} \in \partial(u + s)(G\bar{w})$; the first of these holds by assumption, and to decode the second, note by the convexity of H and F as well as the fact that u and s are maximizations over linear functions that

$$\begin{aligned}\partial u(p) &= \text{conv}(\{-q \in H : (v - p)^\top q = u(p)\}) = -\{q \in H : (v - p)^\top q = u(p)\} = -U(p), \\ \partial s(p) &= \text{conv}(\{q \in F : p^\top q = s(p)\}) = \{q \in F : p^\top q = s(p)\} = S(p),\end{aligned}$$

whereby the rule $\partial(u + s) = \partial u + \partial s$ for finite convex functions and the assumptions $\bar{q}' \in -U(G\bar{w})$ and $\bar{q} + \bar{q}' \in -S(G\bar{w})$ grant

$$-\bar{q} \in S(G\bar{w}) + \{\bar{q}'\} \in S(G\bar{w}) - U(G\bar{w}) = \partial(u + s)(G\bar{w}),$$

which establishes the second optimality condition.

It remains to be shown that \bar{q}' is also optimal. For this, applying first order sufficient conditions to $q' \mapsto u^*(q') + s^*(-\bar{q} - q')$, it suffices to show that

$$0 \in \partial u^*(\bar{q}') + \partial(q' \mapsto s^*(-\bar{q} - q'))(\bar{q}') = \partial u^*(\bar{q}') - \partial s^*(-\bar{q} - \bar{q}'),$$

where the second equality used composition rules for subdifferentials [Hiriart-Urruty and Lemaréchal 2001, Theorem D.4.2.1]. This in turn completes the proof, since standard conjugacy rules [Hiriart-Urruty and Lemaréchal 2001, Proposition E.1.4.3] grant $G\bar{w} \in \partial u^*(\bar{q}')$ via $\bar{q}' \in \partial u(G\bar{w}) = -U(G\bar{w})$ and $G\bar{w} \in \partial s^*(-\bar{q} - \bar{q}')$ via $\bar{q} + \bar{q}' \in -\partial s(G\bar{w}) = -S(G\bar{w})$, whereby $0 = G\bar{w} - G\bar{w} \in \partial u^*(\bar{q}') - \partial s^*(-\bar{q} - \bar{q}')$. \square

LEMMA B.3. *Let closed convex bounded below $r : \mathbf{R}^d \rightarrow \mathbf{R}$, matrix $G \in \mathbf{R}^{n\ell \times d}$ and vector $v \in \mathbf{R}^{n\ell}$ be given. Suppose further that r is either constant or has compact level sets. Then given any probability measure ν over $v \in \mathbf{R}^{n\ell}$ with $\mathbf{E}_\nu[\|v\|_\infty] < \infty$, the function*

$$h(w) = \mathbf{E}_\nu[u(Gw; v) + s(Gw) + r(w)] \tag{19}$$

attains a minimum.

PROOF. If r has compact level sets, then the result follows since the rest of h is bounded below, and thus h itself has compact level sets and attains a minimum. Otherwise, suppose r is equal to some constant $c \in \mathbf{R}$ everywhere; in this case, h is invariant over $\ker(G)$ (due to w only appearing as Gw now that r is constant), so it is convenient to explicitly rule out changes along $\ker(G)$ and consider the auxiliary function

$$f(w) = h(w) + \iota_{\ker(G)^\perp}(w).$$

We will show that f is 0-coercive, and thus has compact level sets and attains a minimum [Hiriart-Urruty and Lemaréchal 2001, Proposition B.3.2.4]. This in turn completes the proof, since a minimum for f is also a minimum of h as follows. For any $w \in \mathbf{R}^d$ consider the decomposition $w = w_\perp + w_k$ where $w_\perp \in \ker(G)^\perp$ is the orthogonal projection of w onto $\ker(G)^\perp$ and $w_k \in \ker(G)$ is the orthogonal projection of w onto $\ker(G)$. Since h is invariant to $\ker(G)$, then $h(w) = h(w_\perp)$. But as h and f agree over $\ker(G)^\perp$, we in fact have $h(w) = h(w_\perp) = f(w_\perp)$, and moreover

$$\inf\{h(w) : w \in \mathbf{R}^d\} = \inf\{h(w) : w \in \ker(G)^\perp\} = \inf\{f(w) : w \in \ker(G)^\perp\}.$$

Since f is infinite off of $\ker(G)^\perp$, its minimum occurs along $\ker(G)^\perp$, and the above equalities grant that this minimum is also a minimum for h .

To prove 0-coercivity of f , let $w \in \mathbf{R}^d$ be an arbitrary nonzero direction, and note that it suffices to consider $w \in \ker(G)^\perp$ (since f is $+\infty$ in other directions). Let $p = Gw$. There are now two cases to consider on the sign of $p_{\min} = \min\{p_i(x) : i \in [n], x \in X\}$:

either $p_{\min} < 0$, or not. If $p_{\min} < 0$, then Lemma B.1 grants

$$\lim_{t \rightarrow \infty} \frac{f(tw) - f(0)}{t} = \lim_{t \rightarrow \infty} \frac{\mathbf{E}_\nu [u(tp; v)] + s(tp) + c - c}{t} \geq \lim_{t \rightarrow \infty} \frac{\mathbf{E}_\nu [-\|v\|_\infty - tp_{\min}]}{t} = -p_{\min} > 0,$$

which means f is 0-coercive [Hiriart-Urruty and Lemaréchal 2001, Proposition B.3.2.4]. For the other case, $w \neq 0$ combined with $w \in \ker(G)^\perp$ implies that $p = Gw \neq 0$. Thus $p_{\min} \geq 0$ means that $p_{\max} = \max\{p_i(x) : i \in [n], x \in X\} > 0$. Once again invoking Lemma B.1,

$$\lim_{t \rightarrow \infty} \frac{f(tw) - f(0)}{t} = \lim_{t \rightarrow \infty} \frac{\mathbf{E}_\nu [u(tp; v)] + s(tp) + c - c}{t} \geq \lim_{t \rightarrow \infty} \frac{tp_{\max}}{t} = p_{\max} > 0,$$

proving 0-coercivity [Hiriart-Urruty and Lemaréchal 2001, Proposition B.3.2.4]. \square

LEMMA B.4. *Consider the setting of Lemma B.3, providing objects G and ν , but suppose $r(w) = \lambda \|w\|_2^2/2$ for some $\lambda \geq 0$. Let h denote the function in (19), let \bar{w} denote any minimizer of h (as provided by Lemma B.3), and set $\mathcal{W}_0 = \{w \in \mathbf{R}^d : h(w) \leq h(\bar{w})\}$.*

—If $G = G_{\text{id}}$, then every $w \in \mathcal{W}_0$ (including \bar{w}) satisfies $\|w\|_\infty \leq (n+1)V$ and $\|w\|_2 \leq (n+1)V\sqrt{\ell}$.

—If $G = G_{p(r)}$ for some integer $r \geq 1$, then every $w \in \mathcal{W}_0$ (including \bar{w}) satisfies $\|w\|_\infty \leq (n+1)V2^r$ and $\|w\|_2 \leq (n+1)Vm^{r/2}2^r$.

PROOF. Before specializing the choice of G , there are a few general properties to note. First we have

$$h(0) = \mathbf{E}_\nu [u(0; v)] + 0 + 0 = \mathbf{E}_\nu \left[\max_{q \in H} v^\top q \right] \leq \mathbf{E}_\nu [n\|v\|_\infty] = nV,$$

where the bound follows from Hölder's inequality. In particular $w \in \mathcal{W}_0$ implies $h(\bar{w}) \leq h(w) \leq h(0) \leq nV$. Consequently, for any $w \in \mathcal{W}_0$, letting $p = Gw$, we have from Lemma B.1 that

$$nV \geq h(w) \geq \mathbf{E}_\nu [u(p; v)] + 0 + 0 \geq -\mathbf{E}_\nu [\|v\|_\infty] - \min_{i \in [n], x \in X} p_i(x)$$

and

$$nV \geq h(w) \geq 0 + s(p) + 0 \geq \max_{i \in [n], x \in X} p_i(x),$$

which together imply

$$\|p\|_\infty = \|Gw\|_\infty \leq (n+1)V. \quad (20)$$

Now consider the case $G = G_{\text{id}}$. In this case, (20) directly provides $\|w\|_\infty \leq (n+1)V$ for $w \in \mathcal{W}_0$, and thus $\|w\|_2 \leq (n+1)V\sqrt{\ell}$.

The remainder of the proof will handle the case $G = G_{p(r)}$ for some integer $r \geq 1$, with an arbitrary $w \in \mathcal{W}_0$ as before. The feature space now has one component for each bundle of size at most r ; thus we write $w(x)$ for the component of w corresponding to $x \in X$ where $|x| \leq r$. Let $x \in X$ be of size at most r . We have

$$p(x) = \sum_{x' \subseteq x} w(x').$$

By Möbius inversion [Bender and Goldman 1975, Theorem 1, Example 2],

$$w(x) = \sum_{x' \subseteq x} (-1)^{|x \setminus x'|} p(x'),$$

and therefore

$$|w(x)| \leq 2^r(n+1)V.$$

Thus $\|w\|_\infty \leq (n+1)V2^r$ and $\|w\|_2 \leq (n+1)V2^r\sqrt{d}$, where d is the dimension of the feature space under $G_{p(r)}$. Applying the simple bound $d \leq m^r$ yields the result. \square

The following convergence result is standard; for other versions, see [Bubeck 2014].

LEMMA B.5. *Let the following objects be given as specified.*

- A closed convex (but not necessarily bounded) constraint set $\mathcal{W} \subseteq \mathbf{R}^d$.
- Iterates $(w^t)_{t \geq 1}$ and supporting objects $(g^t, \eta^t)_{t \geq 1}$ with $w^1 \in \mathcal{W}$ arbitrary, $\eta^t = c/\sqrt{t}$ for some scalar $c > 0$, and $w^{t+1} = \Pi_{\mathcal{W}}(w^t - \eta^t g^t)$ where $g^t \in \mathbf{R}^d$ is arbitrary.
- A sequence of finite convex functions $(f^t)_{t \geq 1}$, where δf^t will denote an arbitrary sub-gradient of f^t at $w^t \in \mathbf{R}^d$.

Then for any time horizon $T \geq 1$, any comparator $u \in \mathcal{W}$, and any $L \geq \sup_{t \in [T]} \|g^t\|_2$,

$$\begin{aligned} & \frac{1}{\sum_{t=1}^T \eta^t} \sum_{t=1}^T \eta^t (f^t(w^t) - f^t(u)) \\ & \leq \frac{1}{c\sqrt{T}} \left(\|w^1 - u\|_2^2 + 2 \sum_{t=1}^T \eta^t (w^t - u)^\top (\delta f^t - g^t) + L^2 c^2 \ln(eT) \right). \end{aligned}$$

PROOF. For any $t \geq 1$, by properties of orthogonal projection onto closed convex sets [Hiriart-Urruty and Lemaréchal 2001, Proposition 3.1.3],

$$\begin{aligned} \|w^{t+1} - u\|_2^2 &= \|\Pi_{\mathcal{W}}(w^t - \eta^t g^t) - \Pi_{\mathcal{W}}(u)\|_2^2 \\ &\leq \|w^t - \eta^t g^t - u\|_2^2 \\ &= \|w^t - u\|_2^2 - 2\eta^t (w^t - u)^\top g^t + (\eta^t)^2 \|g^t\|_2^2 \\ &\leq \|w^t - u\|_2^2 - 2\eta^t (w^t - u)^\top \delta f^t + 2\eta^t (w^t - u)^\top (\delta f^t - g^t) + (\eta^t)^2 L^2, \end{aligned}$$

which by rearrangement and the definition of δf^t gives

$$\begin{aligned} 2\eta^t (f^t(w^t) - f^t(u)) &\leq -2\eta^t (u - w^t)^\top \delta f^t \\ &\leq \|w^t - u\|_2^2 - \|w^{t+1} - u\|_2^2 + 2\eta^t (w^t - u)^\top (\delta f^t - g^t) + (\eta^t)^2 L^2. \end{aligned}$$

Summing across $t \in [T]$,

$$\begin{aligned} & \frac{1}{\sum_{t=1}^T \eta^t} \sum_{t=1}^T \eta^t (f^t(w^t) - f^t(u)) \\ & \leq \frac{1}{2 \sum_{t=1}^T \eta^t} \left(\|w^1 - u\|_2^2 - \|w^{T+1} - u\|_2^2 + 2 \sum_{t=1}^T \eta^t (w^t - u)^\top (\delta f^t - g^t) + L^2 \sum_{t=1}^T (\eta^t)^2 \right), \end{aligned}$$

and the result follows from the estimates

$$\sum_{t=1}^T \eta^t \geq c \int_1^{T+1} \frac{dx}{\sqrt{x}} \geq 2c(\sqrt{T+1} - 1) \quad \text{and} \quad \sum_{t=1}^T (\eta^t)^2 \leq c^2(1 + \ln(T)),$$

as well as the elementary inequality $4(\sqrt{T+1} - 1) \geq \sqrt{T}$. \square

The following result is closely related to Afriat's theorem, for which several proofs are available [Fostel et al. 2004].

LEMMA B.6. *Let $(b^t)_{t \geq 1}$ be the sequence of bids and $(p^t)_{t \geq 1}$ be the sequence of prices. There exists a single value vector v that is consistent with b^t under prices p^t , for all $t \in [T]$, if and only if the sequence of bids satisfies the GARP activity rule (17) with respect to the sequence of prices.*

PROOF. Given the sequences of bids and prices, form the following system of linear inequalities in the variables ζ_i^t for each $i \in [n]$ and $t \in [T]$:

$$\zeta_i^t - p_i^{t\top} b_i^t \geq \zeta_i^s - p_i^{s\top} b_i^s \quad (i \in [n], s, t \in [T]). \quad (21)$$

Feasibility of these inequalities is a necessary condition for there to exist a valuation v consistent with each bid, because they must hold for $\zeta_i^t = v_i^\top b_i^t$. To establish sufficiency, define

$$v_i(x) = \min_{t \in [T]} \{ \zeta_i^t - p_i^{t\top} b_i^t + p_i^t(x) \} \quad (22)$$

for $i \in [n]$ and $x \in X$. If x^t is the bundle associated with b_i^t (recall that bid vectors in the auction are integer), then $v_i^\top b_i^t = v_i(x^t) = \zeta_i^t$ because, by (21),

$$\zeta_i^t - p_i^{t\top} b_i^t + p_i(x^t) = \zeta_i^t \leq \min_{s \in [T]} \{ \zeta_i^s - p_i^{s\top} b_i^s + p_i^s(x^t) \}.$$

Now let e_j for $j = 1, \dots, \ell$ be the unit vectors in \mathbf{R}^ℓ and let e_0 be the origin. Consider any $b_i \in H_i$, which can be written as a convex combination $b_i = \sum_{j=0}^{\ell} \alpha_j e_j$ for non-negative weight α_j that sum to 1. We have

$$(v_i - p_i^t)^\top (b_i^t - b_i) = \sum_{j=0}^{\ell} \alpha_j [(v_i - p_i^t)^\top (b_i^t - e_j)] = \sum_{j=0}^{\ell} \alpha_j [\zeta_i^t - p_i^{t\top} b_i^t + p_i^{t\top} e_j - v_i^\top e_j] \geq 0,$$

where the last inequality follows from (22). Therefore the v defined in (22) is consistent with all bids in the sequence at the given prices, and there exists a single value vector v consistent with all bids if and only if (21) is feasible.

By the Farkas lemma, inequalities (21) are feasible if and only if the optimal value of the following linear program is 0. The LP has a non-negative variable λ_{st} for each $s, t \in [T]$.

$$\begin{aligned} \min_{\lambda \geq 0} \quad & \sum_{s, t \in [T]} (b^t - b^s)^\top p^s \lambda_{st} \\ \text{s.t.} \quad & \sum_{s \in [T]} \lambda_{st} + \sum_{s \in [T]} \lambda_{ts} = 0 \quad (t \in [T]). \end{aligned}$$

This is exactly the LP corresponding to a minimum cost circulation problem over a complete directed graph with a node for each $t \in [T]$, where the cost of edge (s, t) is $(b^t - b^s)^\top p^s$. See [Bertsekas 1998, Chapter 4] for the equivalence of min-cost flow and circulation problems to their LP formulations. If (17) does not hold, then there is a negative cost cycle, and the value of the LP is negative (in fact, unbounded below). Conversely, assume (17) holds. By the conformal realization theorem [Bertsekas 1998, Proposition 1.1], any circulation can be decomposed into a sum of simple cycle flows. As (17) implies that the cost of each simple cycle is non-negative, the cost of the circulation itself is non-negative, by linearity of cost. The optimal value of the LP is therefore 0, which is achieved by setting each λ_{st} to 0. Thus (21) is feasible if and only if the bids satisfy the GARP activity rule. \square

The following result bounds the maximum value, across all agents and bundles, under the standard logit (Gumbel) and probit (Gaussian) random utility models. The result also covers the more general case of any subgaussian error term, and we stress that error components may be correlated in this case.

LEMMA B.7. *Consider the following choices of distribution ν over $v \in \mathbf{R}^{n\ell}$, which takes the form $v = \tilde{v} + \epsilon$. In each case $\tilde{v} \in \mathbf{R}^{n\ell}$ and $\sigma \in \mathbf{R}^{n\ell}$ are deterministic quantities, with $\sigma_{\max} = \max_i \sigma_i$ for convenience, and $\epsilon \in \mathbf{R}^{n\ell}$ has zero mean.*

—If ϵ_i is drawn from a Gumbel distribution with scale parameter σ_i , then

$$\mathbf{E}_\nu[\|v\|_\infty] \leq \|\tilde{v}\|_\infty + 2\sigma_{\max} \ln(2n\ell\sqrt{\pi}).$$

—If ϵ_i is drawn from a Gaussian with mean \tilde{v} and variance σ_i^2 , then

$$\mathbf{E}_\nu[\|v\|_\infty] \leq \|\tilde{v}\|_\infty + \sigma_{\max} \sqrt{2 \ln(2n\ell)}.$$

—If ϵ_i is subgaussian with parameters $(0, \sigma_i^2)$, then

$$\mathbf{E}_\nu[\|v\|_\infty] \leq \|\tilde{v}\|_\infty + \sigma_{\max} \sqrt{2 \ln(2n\ell)}.$$

PROOF. First recall the following standard derivation linking maxima of random variables and their moment generating functions [e.g., Boucheron et al. 2013, Section 2.5]. For any (possibly dependent) random variables (X_1, \dots, X_M) , note by convexity for any $t \geq 0$ that

$$\begin{aligned} \exp\left(t \mathbf{E}_\nu[\max_i |X_i|]\right) &\leq \mathbf{E}_\nu \left[\exp(t \max_i |X_i|) \right] \\ &= \mathbf{E}_\nu \left[\max_i \exp(t|X_i|) \right] \\ &\leq \mathbf{E}_\nu \left[\sum_i \exp(t|X_i|) \right] \\ &\leq \sum_i (\mathbf{E}_\nu[\exp(tX_i)] + \mathbf{E}_\nu[\exp(-tX_i)]). \end{aligned} \quad (23)$$

For the Gumbel distributions with scale parameter σ_i , first note for every $s \in \mathbf{R}$ that $\mathbf{E}_\nu[\exp(s\epsilon_i)] = \Gamma(1 - s\sigma_i)$. Combining this with (23), we have

$$\mathbf{E}_\nu[\max_i |\epsilon_i|] \leq \frac{1}{t} \ln \left(\sum_i (\Gamma(1 - t\sigma_i) + \Gamma(1 + t\sigma_i)) \right),$$

whereby the choice $t = 1/(2\sigma_{\max})$ combined with properties of Γ grants

$$\mathbf{E}_\nu[\max_i |\epsilon_i|] \leq 2\sigma_{\max} \ln \left(\sum_i (\Gamma(1 - \sigma_i/(2\sigma_{\max})) + \Gamma(1 + \sigma_i/(2\sigma_{\max}))) \right) \leq 2\sigma \ln(2n\ell\sqrt{\pi}),$$

and the desired bound follows since $\mathbf{E}_\nu[\|v\|_\infty] \leq \|\tilde{v}\|_\infty + \mathbf{E}_\nu[\|\epsilon\|_\infty]$.

Next consider the bound on arbitrary subgaussian random variables, which will immediately grant the Gaussian bound. As each ϵ_i satisfies $\mathbf{E}_\nu[\exp(t\epsilon_i)] \leq \exp(t^2\sigma_i^2/2)$ by

definition, (23) gives

$$\begin{aligned} \exp\left(t\mathbf{E}_\nu[\max_i |\epsilon_i|]\right) &\leq \sum_i (\mathbf{E}_\nu[\exp(t\epsilon_i)] + \mathbf{E}_\nu[\exp(-t\epsilon_i)]) \\ &\leq 2 \sum_i \exp(t^2\sigma_i^2/2) \\ &\leq 2n\ell \exp(t^2\sigma_{\max}^2/2). \end{aligned}$$

We therefore obtain

$$\mathbf{E}_\nu[\max_i |\epsilon_i|] \leq \frac{\ln(2n\ell)}{t} + \frac{t\sigma_{\max}^2}{2}.$$

Since this expression holds for all $t > 0$, using the optimal choice $t = \sqrt{2\ln(2n\ell)/\sigma_{\max}^2}$ gives the result since $\mathbf{E}_\nu[\|v\|_\infty] \leq \|\tilde{v}\|_\infty + \mathbf{E}_\nu[\|\epsilon\|_\infty]$. \square