

A Bayesian Clearing Mechanism for Combinatorial Auctions

Gianluca Brero
University of Zurich
brero@ifi.uzh.ch

Sébastien Lahaie
Google Research
slahaie@google.com

Abstract

We cast the problem of combinatorial auction design in a Bayesian framework in order to incorporate prior information into the auction process and minimize the number of rounds to convergence. We first develop a generative model of agent valuations and market prices such that clearing prices become maximum a posteriori estimates given observed agent valuations. This generative model then forms the basis of an auction process which alternates between refining estimates of agent valuations and computing candidate clearing prices. We provide an implementation of the auction using assumed density filtering to estimate valuations and expectation maximization to compute prices. An empirical evaluation over a range of valuation domains demonstrates that our Bayesian auction mechanism is highly competitive against the combinatorial clock auction in terms of rounds to convergence, even under the most favorable choices of price increment for this baseline.

Introduction

Combinatorial auctions address the problem of allocating multiple distinct items among agents who may view the items as complements or substitutes. In such auctions, agents can place bids on entire packages of items in order to express complex preferences, leading to higher allocative efficiency. Nevertheless, bidding in a combinatorial auction places a substantial cognitive burden on agents, because the process of valuing even a single bundle can be a costly exercise (Kwasnica et al. 2005; Parkes 2006). There is therefore great interest in developing *iterative* combinatorial auctions, which help to guide the bidding process using price feedback, and in devising techniques to limit the number of rounds needed to reach convergence (ideally in the dozens rather than hundreds) (Petraakis, Ziegler, and Bichler 2012; Bichler, Hao, and Adomavicius 2017).

In this work, we propose to incorporate *prior information* on agent valuations into the auction procedure in a principled manner, thereby achieving a low number of rounds in practice. We cast the problem of combinatorial auction design in a Bayesian framework by developing a joint generative model of agent valuations and market prices. Our generative model defines a likelihood function for clearing prices given

agent valuations. If these valuations are observed, the maximum a posteriori (MAP) estimate for prices corresponds to market clearing prices. If they remain latent, valuations can be marginalized away, weighed by their own likelihood according to observed bids. This forms the basis for an auction scheme to solve the more general clearing problem where valuations are unknown.

We consider settings where several indivisible items are up for sale, and agents have super-additive valuation functions over bundles of items (i.e., the items are pure complements). We provide an auction implementation using item prices consisting of two components. In the *knowledge update* component, we maintain a Gaussian posterior over agent valuations, which is updated as new bids are placed using assumed density filtering (Opper and Winther 1998). Prior information can be incorporated into the auction by suitably initializing this component. The knowledge update step presumes that agents follow *myopic best-response* strategies and bid on utility-maximizing bundles at each round. Accordingly, we discuss an extension to our auction scheme using multiple price trajectories that incentivizes this behavior in *ex post* Nash equilibrium. In the *price update* component, we obtain an analytical expression for the clearing price objective, based on the Gaussian model of valuations that the auction maintains. We establish that the form of the objective is suitable for optimization using expectation maximization. By alternating the two components, we obtain an intuitive and tractable auction scheme where agents place bids, knowledge over latent valuations is updated given bids, and prices are updated given current knowledge of valuations.

For evaluation purposes, we first illustrate our auction on a stylized instance to gain insight into the auction’s behavior under both unbiased and biased prior information. We then conduct simulation experiments to compare our auction implementation against a combinatorial clock auction that updates prices according to excess demand, which is the standard price update scheme used in practice (Ausubel and Baranov 2014). The prior information in our Bayesian auction is obtained by fitting a Gaussian process prior on a training sample of valuations. The baseline clock auction is parametrized by a step size, or price increment. We find in our experiments that our Bayesian auction is competitive against the strongest possible version of the baseline auction, where the price increment is chosen separately for each in-

stance to lead to the fewest possible rounds. In particular, the Bayesian auction almost matches the strongest possible version of baseline auction in terms of number of instances cleared, and uses fewer rounds on average when it is able to clear.

Preliminaries

We consider a setting with m distinct and indivisible items, held by a single seller. The items are to be allocated among n agents (i.e., buyers). We will use the notation $[n] = \{1, \dots, n\}$, so that $[n]$ and $[m]$ denote the index sets of agents and items, respectively. There is unit supply of each item. A *bundle* is a subset of the set of items. We associate each bundle with its indicator vector, and denote the set of bundles as $\mathcal{X} = \{0, 1\}^m$. The component-wise inequality $x \leq x'$ therefore means that bundle x is contained in bundle x' . The empty bundle is denoted by \emptyset .

Each agent i is *single-minded* so that its valuation can be encoded via a pair (x_i, w_i) where $x_i \in \mathcal{X}$ is a bundle and $w_i \in \mathbb{R}_+$ is a non-negative value (i.e., willingness to pay) for the bundle. The agent's valuation function $v_i : \mathcal{X} \rightarrow \mathbb{R}_+$ is defined as $v_i(x) = w_i$ if $x \geq x_i$, and $v_i(x) = 0$ otherwise. In words, the agent only derives positive value if it acquires all the items in x_i (which are therefore complements), and any further item is superfluous. Our auction and results all extend to agents with OR valuations, which are concise representations of super-additive valuations (Nisan 2000).¹ This is due to the fact that an agent with an OR valuation will behave and bid in our auction exactly like a set of single-minded agents, under myopic best-response (Parkes 1999). Under super-additive valuations, items are pure complements, and complementarities are a key motivation for using package bidding. For the sake of simplicity, however, we limit the exposition to single-minded agents.

An *allocation* is represented as a vector of bundles $\mathbf{y} = (y_1, \dots, y_n)$, listing the bundle that each agent obtains (possibly \emptyset). An allocation is *feasible* if the listed bundles are pairwise disjoint (i.e., each item is allocated to at most one agent). We denote the set of feasible allocations by \mathcal{F} . The purpose of running a combinatorial auction is to find an *efficient* allocation of the items to the agents, meaning an allocation that maximizes the total value to the agents.² More formally, a feasible allocation $\mathbf{y} \in \mathcal{F}$ is efficient if $\mathbf{y} \in \arg \max_{\mathbf{y}' \in \mathcal{F}} \sum_{i \in [n]} v_i(y'_i)$. However, an iterative auction proceeds via a price adjustment process, so prices will be our central object of study, rather than allocations. The allocation in an iterative auction is adjusted according to agents' responses to prices.

¹More formally, an OR valuation takes the form $v(x) = \max\{v_1(x_1) + v_2(x_2) : x_1 + x_2 = x, x_1 \wedge x_2 = 0\}$, where v_1 and v_2 are themselves OR valuations or single-minded.

²This is in contrast to the goal of maximizing *revenue*. In auction design, one typically begins with an efficient auction, which is then modified (e.g., using reserve prices) to achieve optimal revenue (Myerson 1981). We therefore consider the problem of designing an efficient auction as more fundamental.

Clearing Prices In the context of a combinatorial auction, we encode prices as a non-negative function $\theta : \mathcal{X} \rightarrow \mathbb{R}_+$ over the bundles. We assume that prices are normalized and monotone: $\theta(\emptyset) = 0$, and $\theta(x) \leq \theta(x')$ if $x \leq x'$. An iterative auction adjusts prices to balance demand and supply. To formalize this notion, we need several additional concepts. We assume that agents have quasi-linear utility, so that the utility to agent i of obtaining bundle x at prices θ is $v_i(x) - \theta(x)$. The *indirect utility* function provides the maximum utility that agent i can achieve, when faced with prices θ , by choosing among bundles from \mathcal{X} :

$$V_i(\theta) = \max \{ v_i(x) - \theta(x) : x \in \mathcal{X} \}. \quad (1)$$

Note that for single-minded agents, the indirect utility reduces to $V_i(\theta; x_i, w_i) = (w_i - \theta(x_i))_+$, where the notation $(a)_+ = \max\{a, 0\}$ refers to the positive part of the argument. It will sometimes be useful to make explicit the parametrization of the indirect utility on the agent's type (x_i, w_i) , as we have just done. The *demand set* of agent i is defined as $D_i(\theta) = \{x \in \mathcal{X} : v_i(x) - \theta(x) = V_i(\theta)\}$. Similarly, the *indirect revenue* function provides the maximum revenue that the seller can achieve, when faced with prices θ , by selecting among feasible allocations:

$$R(\theta) = \max \left\{ \sum_{i \in [n]} \theta(y_i) : \mathbf{y} \in \mathcal{F} \right\}. \quad (2)$$

The seller's *supply set* consists of the feasible allocations that maximize revenue: $S(\theta) = \{\mathbf{y} \in \mathcal{F} : \sum_{i \in [n]} \theta(y_i) = R(\theta)\}$.

We say that prices θ are *clearing prices* if there is a feasible allocation \mathbf{y} such that, at prices θ , the seller's revenue is maximized, and each agent's utility is maximized. Formally, we require the following conditions: $\mathbf{y} \in S(\theta)$ and $y_i \in D_i(\theta)$ for all $i \in [n]$. We say that the clearing prices θ *support* allocation \mathbf{y} .

It is a standard result that the set of allocations supported by any given clearing prices θ coincides with the set of efficient allocations. (This is a special case of the Fundamental Theorems of Welfare Economics (Mas-Colell, Whinston, and Green 1995, 16.C–D).) Moreover, Bikhchandani and Ostroy (2002) have shown that clearing prices exist and coincide with the minimizers of the following objective function, which corresponds to the linear programming dual of the problem of allocating the items efficiently:

$$\sum_{i \in [n]} V_i(\theta) + R(\theta). \quad (3)$$

This is a piece-wise linear, convex function of the price function θ . Importantly, this result is guaranteed to hold only if the prices are an unrestricted function over the bundles (except for non-negativity and normalization). In practice, it is common to use certain parametrizations for the prices. For instance, taking $\theta(x) = p \cdot x$ for some vector $p \in \mathbb{R}_+^m$ corresponds to using linear prices (i.e., item prices). These parametrizations may not achieve the unrestricted minimum in (3); in particular, linear clearing prices may not exist. We will use unrestricted prices in the development of our

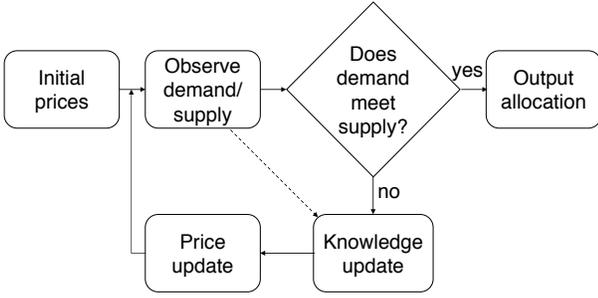


Figure 1: Bayesian iterative auction.

auction, and postpone the question of price parametrization until needed to achieve a practical implementation.

It is useful to view (3) as a potential function that quantifies how close prices θ are to supporting an efficient allocation. Indeed, if some prices achieve a value of (3) that differs from the optimum by an additive error of δ , then the agents (and seller) can be induced to accept an efficient trade using transfers totaling δ . In the following, we will therefore refer to the function

$$U(\theta; \mathbf{v}) = \exp \left[- \sum_{i \in [n]} V_i(\theta) - R(\theta) \right] \quad (4)$$

as the *clearing potential* for the valuation profile $\mathbf{v} = (v_1, \dots, v_n)$, which will capture, in a formal sense, how likely a price function θ is to clearing valuation profile \mathbf{v} .

Iterative Auction and Incentives The goal of our paper is to design an iterative auction that exploits the auctioneer’s prior knowledge over agent valuations in order to speed up the clearing process. The auction proceeds over rounds. Agents report their demand at the current prices and, if the market is not cleared, the information provided by the agents is used to update the knowledge about their valuations. Candidate clearing prices are computed based on the updated knowledge, and the procedure iterates. A schematic representation of the auction process is presented in Figure 1. The knowledge update and price update components constitute the core of the auction that must be implemented.

The correctness of our auction relies on the agents following a strategy of *myopic best-response* bidding, meaning that each agent bids on a utility-maximizing bundle at each round. There is evidence that myopic bidding may be a reasonable assumption in practice. For instance, in the FCC broadband spectrum auction, jump bids were the exception (Cramton 1997). Nonetheless, a robust auction design should incentivize agents to follow the appropriate strategies. For this purpose, we can use an extension of our auction that maintains $n + 1$ price trajectories in order to compute clearing prices when all agents are present, and when each agent is removed in turn. This allows one to compute final VCG payments and bring myopic best-response bidding into an *ex post* Nash equilibrium (Gul and Stacchetti 2000;

Bikhchandani and Ostroy 2006). The technique of using multiple trajectories was previously used by Ausubel (2006) and Mishra and Parkes (2007) among others. We will provide a more precise treatment of incentives in the formal description of our auction mechanism.

Generative Model

The purpose of this section is to define a probabilistic relationship between prices and valuations that will allow us to use the auctioneer’s prior knowledge over valuations to make inferences over clearing prices. We write $\mathbf{w} = (w_1, \dots, w_n)$ and $\mathbf{x} = (x_1, \dots, x_n)$ for the vectors of agents’ values and bundles, and denote the probabilistic model as $P(\mathbf{x}, \mathbf{w}, \theta)$. Below, our convention is that Q refers to distributions—possibly unnormalized—that form the building blocks of the generative model, whereas P refers to the normalized distribution resulting from the generative model. We represent the prior knowledge of the auctioneer over agent valuations via the probability density function $Q(\mathbf{w}) = \prod_{i \in [n]} Q(w_i)$.

The structure of our probability model is inspired by the work of Sollich (2002), who provides a Bayesian interpretation of the support vector machine (SVM) objective. To establish a proper relationship between prices and valuations, the key is to require that

$$P(\mathbf{x}, \mathbf{w}, \theta) \propto U(\theta; \mathbf{x}, \mathbf{w})Q(\mathbf{w}), \quad (5)$$

where $U(\theta; \mathbf{x}, \mathbf{w})$ is the clearing potential introduced in (4), adapted to single-minded valuations. Under this joint probability model, we have that the posterior probability of prices θ takes the form

$$P(\theta | \mathbf{x}, \mathbf{w}) \propto U(\theta; \mathbf{x}, \mathbf{w}). \quad (6)$$

Therefore, the maximum a posteriori (MAP) estimate maximizes $U(\theta; \mathbf{x}, \mathbf{w})$, or equivalently minimizes (3), and corresponds to clearing prices.

To establish that a probability model of the form (5) is possible—namely, that it can indeed be normalized—we will derive it as the result of a generative model. This process may be of independent interest as a means of generating agents together with market prices.

1. Draw prices θ according to $Q(\theta) \propto \exp[-R(\theta)]$.
2. For each agent $i \in [n]$:
 - Draw $w_i \in \mathbb{R}_+$ from $Q(w_i)$.
 - Draw $x_i \in \mathcal{X}$ from

$$Q(x_i | w_i, \theta) = \frac{1}{2^m} \exp[-V_i(\theta; x_i, w_i)].$$

- With probability $1 - \nu(w_i, \theta)$, restart from step 1, where

$$\nu(w, \theta) = \sum_{x \in \mathcal{X}} Q(x | w, \theta).$$

Above, we must ensure that the $Q(\theta)$ prior normalizes; this is the case under our assumption that the domain of θ falls within the positive orthant. The prior distribution $Q(w_i)$ on value w_i is left free in the model, so that it may correspond to the auctioneer’s prior in practice. Note that the

bundle likelihood $Q(x_i | w_i, \theta)$ is *not* normalized; because $V_i(\theta; x_i, w_i) \geq 0$, summing over the set of bundles leads to the aggregate probability mass $\nu(w_i, \theta) \leq 1$. Rather than normalizing by this quantity, we use the “remaining probability” $1 - \nu(w_i, \theta)$ of not drawing any bundle to restart the process. Because of the possible restart, the agent types (bundle-value pairs) and clearing prices are not independent in the overall generative distribution. In particular, the number of agents n in the economy affects the distribution of prices.

The following proposition confirms that our model satisfies (5). Due to space constraints, all proofs are deferred to the full version of this paper.

Proposition 1. *The generative model of agent types and prices takes the form*

$$P(\mathbf{x}, \mathbf{w}, \theta) \propto U(\theta; \mathbf{x}, \mathbf{w})Q(\mathbf{w}). \quad (7)$$

The generative process defines a probability distribution over prices once valuations are *observed*, but during the auction the valuations remain latent, and must be inferred based on observed bids placed across rounds. Under appropriate incentives, the auctioneer can infer valuations assuming that the agents follow myopic best-response bidding. However, if there are any bidding errors or corruption in communication, assuming exact best-response can cause singularities in the inference process (e.g., there may be no valuation consistent with all observed bids). To guard against this, our mechanism will integrate bids as if they were generated from the following stochastic model: Let $b_i \in \{-1, +1\}$ be an indicator variable to denote whether the agent bids on bundle x_i ($b_i = +1$) or not ($b_i = -1$); the latter is equivalent to bidding on \emptyset . If the cost of bundle x_i is c_i , then the choice of bid follows the probability distribution

$$Q(b_i = +1 | c_i, w_i) \propto \Phi(\beta(w_i - c_i)), \quad (8)$$

where Φ is the cumulative distribution function of the standard normal, and $\beta > 0$ is a scalar parameter. This is known as the *probit* variant of approximate best-response, which arises from random utility models (Train 2009). As $\beta \rightarrow \infty$, we obtain exact best-response: the agent bids on x_i if and only if this bundle yields positive utility under bundle cost c_i . Using a large (but finite) β allows the auctioneer to model agents as essentially following a best-response strategy, but occasionally allowing for bidding errors or inconsistencies.

Auction Description

Our auction proceeds over rounds; we use k to denote the current round, and ℓ to index the rounds up to k . At each round, prices are updated, which imputes a cost to each agent’s bundle. Let c_i^ℓ be the cost of agent i ’s desired bundle x_i in round ℓ according to the current prices. The prices at each round should not be confused with the latent clearing prices θ , which we are trying to compute as a MAP estimate of the generative model. Given its value w_i and the bundle cost c_i^ℓ , agent i places bid $b_i^\ell \in \{-1, +1\}$ in round ℓ .

We write $\mathbf{c}^\ell = (c_1^\ell, \dots, c_n^\ell)$ to denote the vector of bundle costs in round ℓ , and $\mathbf{c}^{(\ell)} = (\mathbf{c}^1, \dots, \mathbf{c}^\ell)$ to denote the vector of costs up to round ℓ . For brevity we also write $\mathbf{c} = \mathbf{c}^{(k)}$ to denote the vector of all costs up to the current round. We use

the notation \mathbf{b}^ℓ , $\mathbf{b}^{(\ell)}$, and \mathbf{b} to denote the analogous vectors of bids. The bundle costs and agent bids in a round depend on the current prices, which themselves depend on the bids placed in all earlier rounds. Assuming that the first round prices are zero, we have the following intuitive posterior over bids and costs.

Lemma 1. *The posterior distribution over bids and costs placed during the auction, given the generated prices and agent types, is given by*

$$P(\mathbf{b}, \mathbf{c} | \mathbf{x}, \mathbf{w}, \theta) = Q(\mathbf{b} | \mathbf{c}, \mathbf{w}) \prod_{\ell=1}^k P(\mathbf{c}^\ell | \mathbf{b}^{(\ell-1)}, \mathbf{c}^{(\ell-1)}, \mathbf{x}),$$

where \mathbf{b}^ℓ and \mathbf{c}^ℓ are the vectors of agent bids and costs at round ℓ , and $\mathbf{b}^{(\ell)}$ and $\mathbf{c}^{(\ell)}$ are the vectors of agent bids and costs up to round ℓ .

We see that the posterior over bids and costs does not depend on the underlying clearing prices θ , conditional on agent types (\mathbf{x}, \mathbf{w}) , because the initial prices and agent valuations fully determine how the auction proceeds. More specifically, the posterior decomposes into the likelihood of the observed bids $Q(\mathbf{b} | \mathbf{c}, \mathbf{w})$ under stochastic model (8), times the likelihood of the observed sequence of costs. The latter does not involve \mathbf{w} , because current round prices are fully determined by the bids and costs of previous rounds. Our auction is based on the following characterization of the overall posterior over prices and agent values.

Proposition 2. *The posterior distribution of latent variables (\mathbf{w}, θ) given observed variables $(\mathbf{b}, \mathbf{c}, \mathbf{x})$ takes the form*

$$P(\mathbf{w}, \theta | \mathbf{b}, \mathbf{c}, \mathbf{x}) \propto Q(\mathbf{b} | \mathbf{c}, \mathbf{w})Q(\mathbf{w}) \cdot Q(\mathbf{x} | \mathbf{w}, \theta)Q(\theta),$$

where the proportionality constant depends solely on $(\mathbf{b}, \mathbf{c}, \mathbf{x})$.

The posterior factors into two terms, which motivates our auction procedure. The first term, $Q(\mathbf{b} | \mathbf{c}, \mathbf{w})Q(\mathbf{w})$, can be construed as a posterior over agent values given bids and costs, since $Q(\mathbf{w})$ corresponds to a prior and $Q(\mathbf{b} | \mathbf{c}, \mathbf{w})$ corresponds to a likelihood. We will maintain an approximation $\hat{P}(\mathbf{w})$ to this posterior over agent values and update it as new bids are placed in response to bundle costs. This is the *knowledge update* component.

Recalling (6), the second term $Q(\mathbf{x} | \mathbf{w}, \theta)Q(\theta) \propto U(\theta; \mathbf{x}, \mathbf{w})$ in the posterior corresponds (up to a constant factor) to the price posterior given knowledge of agent types. This leads to an approximation to the price posterior when values remain latent:

$$P(\theta | \mathbf{b}, \mathbf{c}, \mathbf{x}) \approx \int d\mathbf{w} \hat{P}(\mathbf{w}) \cdot Q(\mathbf{x} | \mathbf{w}, \theta)Q(\theta). \quad (9)$$

Here we have simply integrated the full posterior as given by Proposition 2, and made use of our approximation to the value posterior. (We have also omitted the normalization constant.) In the context of an auction, we quote a specific price function to the agents, rather than a distribution over prices. Therefore, in the *price update* component, we will compute and quote the MAP estimate of prices by maximizing (9). Note that if we have exact knowledge of agent values (i.e., \hat{P} is a point mass), computing the MAP estimate is equivalent to minimizing (3) and to computing clearing prices, as one would expect.

Knowledge Update We observe that the value posterior $Q(\mathbf{b}|\mathbf{c}, \mathbf{w})Q(\mathbf{w})$ consists of a separate factor for each agent i , taking the form $Q(w_i) \prod_{\ell=1}^k Q(b_i^\ell | c_i^\ell, w_i)$, where k is the current round. This represents a posterior on agent i 's individual value w_i . To obtain an approximation to this posterior, we use an online scheme known as *assumed density filtering*, which is a special case of expectation propagation (Cowell, Dawid, and Sebastiani 1996; Minka 2001; Opper and Winther 1998). Under this approach, a Gaussian distribution $\hat{P}(w_i; m_i, \sigma_i^2)$ is used to approximate the posterior; its mean m_i and variance σ_i^2 are updated at each round given the bidding observations. The Gaussian is initially set to approximate the prior $Q(w_i)$ via moment matching: m_i and σ_i^2 are set to the mean and variance of this prior Q . In each later round $\ell = 1, \dots, k$ the posterior is again updated by matching the moments of $Q(b_i^\ell | c_i^\ell, w_i)\hat{P}(w_i; m_i, \sigma_i^2)$, which is an online update. Using moment matching as an approximation is justified by the fact that it corresponds to minimizing the Kullback-Leibler divergence $\text{KL}(Q\|\hat{P})$ under the constraint that \hat{P} is Gaussian.

Due to the form of the likelihood (8) and the fact that \hat{P} is Gaussian, the update has a closed-form solution (see Williams and Rasmussen 2006, p. 74):

$$\begin{aligned} m_i &\leftarrow m_i + \frac{b_i^\ell \sigma_i^2 \beta \mathcal{N}(z_i)}{\sqrt{1 + \sigma_i^2 \beta^2} \Phi(z_i)} \\ \sigma_i^2 &\leftarrow \sigma_i^2 - \frac{\sigma_i^4 \beta^2 \mathcal{N}(z_i)}{1 + \sigma_i^2 \beta^2 \Phi(z_i)} \left(z_i + \frac{\mathcal{N}(z_i)}{\Phi(z_i)} \right), \end{aligned}$$

where \mathcal{N} and Φ are the probability density and cumulative distribution functions of the standard normal, respectively, and where $z_i = b_i^\ell \beta (m_i - c_i^\ell) / \sqrt{1 + \sigma_i^2 \beta^2}$. Recall that β is a positive parameter characterizing the extent to which the auctioneer assumes that agents make mistakes in placing best-response bids. Since β is positive, the mean m_i is updated in the direction of the bid b_i^ℓ . On the other hand, the variance σ_i^2 is strictly decreasing, thus ensuring that the beliefs over bidder values converge to a point mass in the limit as the rounds progress, and that the auction converges to a final vector of prices.

Price Update To implement the price update component we need an algorithm to maximize the approximate posterior (9). This posterior factors into $Q(\theta)$ and a term for each agent i , which we denote as

$$\mathcal{L}_i(\theta) = \int dw_i \hat{P}_i(w_i; m_i, \sigma_i^2) Q(x_i | w_i, \theta).$$

Because $Q_i(x_i | w_i, \theta)$ has an exponential form, and \hat{P}_i is a Gaussian, this integral has a closed form solution. Let $q_i \in \{0, 1\}$ be a binary auxiliary variable. We have $\mathcal{L}_i(\theta) = \sum_{q_i=0}^1 \mathcal{L}_i(\theta, q_i)$, where we define

$$\begin{aligned} \mathcal{L}_i(\theta, 1) &= \Phi \left(\frac{m_i - \theta(x_i)}{\sigma_i} - \sigma_i \right) \exp \left[\theta(x_i) - m_i + \frac{\sigma_i^2}{2} \right] \\ \mathcal{L}_i(\theta, 0) &= \Phi \left(\frac{\theta(x_i) - m_i}{\sigma_i} \right) \end{aligned}$$

Here Φ is again the cumulative distribution function of the standard normal. To summarize, taking the log of (9), the objective we seek to maximize with respect to θ is

$$\log Q(\theta) + \sum_{i \in [n]} \log \sum_{q_i=0}^1 \mathcal{L}_i(\theta, q_i). \quad (10)$$

Now, because Φ is log-concave, both $\mathcal{L}_i(\theta, 0)$ and $\mathcal{L}_i(\theta, 1)$ are log-concave in θ . Ignoring the first term for an instant, we see that the objective consists of a sum of mixtures $\mathcal{L}_i(\theta)$ of log-concave functions for each agent. This kind of objective is well-suited to optimization using the expectation-maximization (EM) algorithm (Dempster, Laird, and Rubin 1977). The q_i amount to ‘‘latent’’ variables and the ‘‘marginal’’ likelihood appears within the objective (10). (However, we do not claim any intuitive interpretation for the latent q_i —they are simply used to fit the objective into the mold of EM.)

The remaining term is $\log Q(\theta)$, which is $-R(\theta)$ up to an additive constant. Recalling the definition of the seller's indirect utility (2), we see that this term is very complex for unrestricted θ , because the set of feasible allocations \mathcal{F} has a very complicated structure. To address this we will impose a linear structure on prices: $\theta(x) = p \cdot x$ where $p \in \mathbb{R}_+^m$ denote item prices. With this parametrization, we have $R(\theta) = \sum_{j=1}^m p_j$, because any allocation that allocates all the items maximizes revenue under linear prices. The $\log Q(\theta)$ term therefore becomes a linear term in p , which is straightforward to incorporate within the EM algorithm.

Incentive Compatibility Our auction converges to an efficient allocation and clearing prices under myopic best-response bidding, but to ensure that agents follow such a strategy, they must be incentivized to do so. The standard technique used to achieve this in the literature on iterative auctions is to charge VCG payments upon completion (Gul and Stacchetti 2000; Bikhchandani and Ostroy 2006). But whereas VCG payments (together with an efficient allocation) induce truthful bidding in dominant strategies for single-shot auctions, weaker results hold for iterative auctions.

A strategy profile constitutes an *ex post* Nash equilibrium if no agent would like to deviate from its strategy, holding the others' strategies fixed, even with knowledge of the private valuations of the other agents. Gul and Stacchetti (2000) prove the following result:

Theorem 1 (Gul and Stacchetti, 2000). *Truthful myopic best-response bidding is an ex post Nash equilibrium in an iterative auction that myopically-implements the VCG outcome.*

Above, the VCG outcome refers to an efficient allocation along with VCG payments, and an auction myopically-implements this outcome if the auction converges to it under myopic best-response bidding. The reason that truthfulness only holds in *ex post* Nash equilibrium, rather than dominant strategies, is that profitable deviations may exist if another agent bids in a manner inconsistent with any valuation.

Our auction already computes the efficient allocation under these conditions by virtue of converging to clearing prices. To compute VCG payments, we can simply extend our auction drawing on the idea of multiple price trajectories: the

usual trajectory traced by our auction, and the trajectories that would result if each agent were removed in turn. This technique was previously used by Ausubel (2006) and Mishra and Parkes (2007). In this extended design, at each round, agents place bids against $n + 1$ different price vectors. Upon completion, the agents place last-and-final bids for their allocated bundles, thereby communicating their value for the allocations; importantly, agents do not need to communicate values for any bundles they did not win. This information is precisely what is needed to compute VCG payments (see, e.g., Parkes and Ungar 2000).

Empirical Evaluation

In this section we evaluate our Bayesian auction design with two different kinds of experiments: a small experiment to illustrate the behavior of our auction under biased and unbiased prior information, and a larger-scale experiment to compare our auction against a competitive baseline.

Our simulations are conducted in Matlab. In all our experiments, we assume that agents best respond to the proposed prices (i.e., they always bid on their most profitable bundle), and that the auctioneer considers their bids as if they were generated from the response model presented in (8) with $\beta = 10$. However, simulations where real bids follows (8) with $\beta = 4$ provide results similar to the ones presented. For the price update, the objective (10) is maximized using the “active-set” algorithm in Matlab. To avoid numerical singularities we place a lower bound of 0.01 on the variance of valuation estimates.

LLG Experiments We consider an instance of the *Local-Local-Global* (LLG) domain (Ausubel and Milgrom 2006), which has been considered several times in the combinatorial auctions literature. There are two items and three single-minded agents. Two of the agents are *local*, meaning that they are interested in just one item, respectively the first and second item. The last agent is *global* in the sense that it is interested in both items.

The two local agents have a value of 4 for their respective items, and the global one has a value of 10 for both. The items are efficiently allocated when they are both assigned to the global agent, and linear prices are expressive enough to clear the market (e.g., we can use a price of 4 for each item).

We assume that the auctioneer has accurate knowledge of the local agents’ values: $Q(w_1) = \mathcal{N}(w_1; m_1 = 4, \sigma_1^2 = 0.01)$ and $Q(w_2) = \mathcal{N}(w_2; m_2 = 4, \sigma_2^2 = 0.01)$. Note the very low variance, reflecting certainty. We test how different kinds of prior knowledge over the global agent’s value affect the number of rounds that the Bayesian auction takes to clear the market. In the first case we assume unbiased prior knowledge: $m_3 = 10$. In the second case we assume that it is biased below: $m_3 = 4$. Here, the auctioneer tends to allocate to the local agents instead of the global one.

Figure 2 plots the number of rounds that our Bayesian auction takes to clear the market against the standard deviation σ_3 of the prior over the global agent’s value. We see that, in the unbiased scenario, the number of rounds monotonically increases as the standard deviation grows. This can be easily

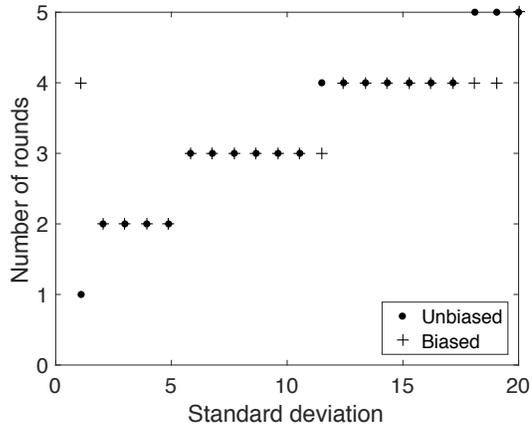


Figure 2: Auction rounds in LLG.

explained since increasing the standard deviation only adds noise to the exact prior estimate. In the biased scenario, we have an optimal range of standard deviations between 2 and 5. If the standard deviation is too low, the auction needs many observations to correct the biased prior. If it is too high, the low confidence leads to many rounds because the auctioneer needs to refine its estimate of the value regardless of the bias.

CATS Experiments For our second set of experiments, we generate instances using the Combinatorial Auction Test Suite (CATS), which offers four generator distributions: *paths*, *regions*, *arbitrary*, and *scheduling* (Leyton-Brown, Pearson, and Shoham 2000). These are meant to model realistic domains such as truck routes, real estate, and pollution rights. We generate 1000 instances from each distribution, each with 12 items and 10 single-minded agents.

The instances are generated as follows. First, 100 input files with 1000 bids each are generated. Each input file is partitioned into a “training set” and “test set”, each with 500 bids. From the test set, 10 bids (representing 10 single-minded agents) are sampled uniformly at random. The training set is used to fit the prior knowledge of our Bayesian auction. Specifically, we fit a linear regression model of bundle value according to items contained, using a Gaussian process with a linear covariance function, leading to Gaussian prior knowledge. The fit is performed using the publicly available GPML Matlab code (Williams and Rasmussen 2006).

As a baseline we implemented a standard linear-price auction scheme closely related to the combinatorial clock auction (Ausubel and Baranov 2014). The scheme is parametrized by a positive step size τ . At each round ℓ , the price of an item is incremented by its *excess demand*, scaled by $\tau/\sqrt{\ell}$. The excess demand of an item is the number of bidded bundles that contain it, minus the number of units of the item offered by the seller at current prices. This can be viewed as a subgradient descent method for computing clearing prices, for which a step size proportional to $1/\sqrt{\ell}$ yields the optimal worst-case convergence rate (Bertsekas

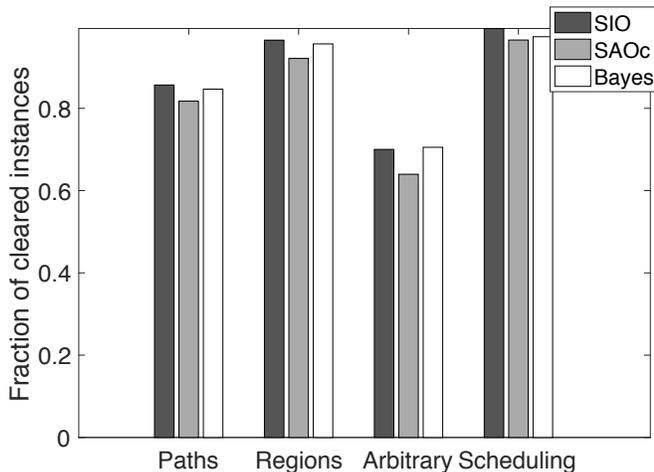


Figure 3: Cleared instances in CATS.

2015, Chap. 3).

Both the Bayesian auction and the baseline subgradient auction use linear prices, but these may not be expressive enough to support an efficient allocation in instances generated by CATS. We therefore set a limit of 100 rounds for each auction run, and record reaching this limit as a failure to clear the market.

On each instance we run a single Bayesian auction, and 100 subgradient auctions with the step size uniformly spanning the interval from zero to the maximum agent value. This leads to several baseline results. The *standard instance optimized* (SIO) results refer to the performance of the baseline when using the optimal step size for each instance. The *standard average clearing-optimized* (SAOc) results refer to the performance of the baseline under the fixed step size that leads to the best clearing performance on average, for each valuation domain. Analogously, the *standard average round-optimized* (SAOr) results refer to baseline performance under the step size leading to lowest average number of rounds. For each instance, the step size that leads to the lowest number of rounds naturally leads to the best clearing rate. But the fixed step sizes that optimize these two criteria on average may be different. Note that SIO is an extremely competitive baseline, since it is optimized for each instance; a priori, we hoped to be competitive against it, but did not expect to beat it. The SAOc and SAOr baselines reflect more realistic performance that could be achieved in practice.

We first consider clearing performance. The results are reported in Figure 3. We find that the Bayesian auction is competitive with SIO on all four domains, and that it always outperforms SAOc. In fact, our auction even outperforms SIO on the *arbitrary* domain. This means that it was able to clear some instances that the subgradient auction could not clear within 100 rounds at any step size. In general, there is good agreement between our Bayesian auction and the baselines on which instances can be cleared or not according to the 100-round criterion. This indicates that failure to clear is typically a property of the instance rather than the algorithm.

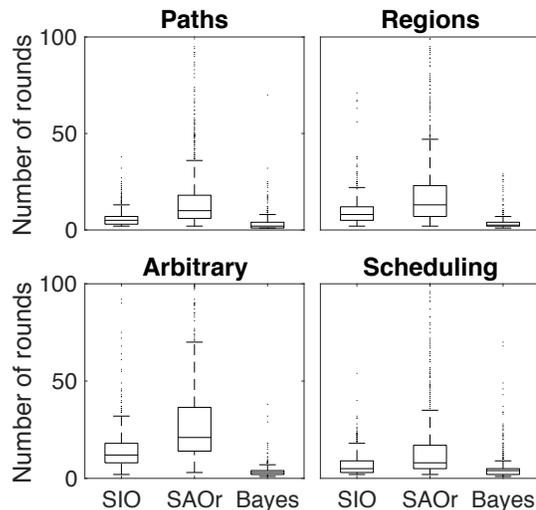


Figure 4: Auction rounds in CATS.

Figure 4 summarizes the distributions of rounds needed to achieve clearing using box plots. To enable fair comparisons, for this plot we only considered instances that were cleared by all auction types: the Bayesian auction, SAOr, and SIO. This yields 770 valid instances for *paths*, 910 for *regions*, 624 for *arbitrary* and 955 for *scheduling*. The mean rounds for the Bayesian auction, SAOr, and SIO are always statistically different at the 0.01 level. We see from the plot that, in terms of the median number of rounds, the Bayesian auction clearly outperforms SAOr, but also remarkably outperforms SIO. Furthermore, the distribution of rounds for the Bayesian auction has a much lower spread than the baselines. It is able to clear almost all instances in less than 10 rounds.

Conclusion

In this work we developed a Bayesian clearing mechanism for combinatorial auctions that allows one to incorporate prior information into the auction process in a principled manner. Our auction mechanism is based on a joint generative model of valuations and prices such that clearing prices are the MAP estimate given observed valuations. Our empirical evaluation confirmed that our Bayesian mechanism performs remarkably well against a conventional clock auction scheme, in terms of rounds to convergence. Our auction’s performance simply relies on reasonable priors for valuations, rather than careful tuning of price increments.

We believe that the Bayesian perspective on auction design developed in this paper could be leveraged to improve other aspects beyond rounds to convergence. For instance, the Bayesian paradigm offers a principled way to select hyperparameters (MacKay 1992); in our context, this could be used to choose the right structure of prices (linear, nonlinear) to clear the market, a priori. The knowledge update component could also form the basis of more refined activity rules; for instance, one could reject bids that are highly unlikely, given the valuation posterior based on previous bids. We intend to pursue these directions in future work.

References

- Ausubel, L. M., and Baranov, O. V. 2014. A practical guide to the combinatorial clock auction. Technical report, University of Maryland.
- Ausubel, L. M., and Milgrom, P. 2006. The lovely but lonely vickrey auction. In Cramton, P.; Shoham, Y.; and Steinberg, R., eds., *Combinatorial auctions*. The MIT Press.
- Ausubel, L. M. 2006. An efficient dynamic auction for heterogeneous commodities. *The American Economic Review* 96(3):602–629.
- Bertsekas, D. P. 2015. *Convex optimization algorithms*. Athena Scientific.
- Bichler, M.; Hao, Z.; and Adomavicius, G. 2017. Coalition-based pricing in ascending combinatorial auctions. *Information Systems Research*. Forthcoming.
- Bikhchandani, S., and Ostroy, J. M. 2002. The package assignment model. *Journal of Economic Theory* 107(2):377–406.
- Bikhchandani, S., and Ostroy, J. M. 2006. Ascending price vickrey auctions. *Games and Economic Behavior* 55(2):215–241.
- Cowell, R. G.; Dawid, A. P.; and Sebastiani, P. 1996. A comparison of sequential learning methods for incomplete data. *Bayesian Statistics* 5:533–542.
- Cramton, P. 1997. The FCC spectrum auctions: An early assessment. *Journal of Economics & Management Strategy* 6(3):431–495.
- Dempster, A. P.; Laird, N. M.; and Rubin, D. B. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B* 1–38.
- Gul, F., and Stacchetti, E. 2000. The english auction with differentiated commodities. *Journal of Economic Theory* 92(1):66–95.
- Kwasnica, A. M.; Ledyard, J. O.; Porter, D.; and DeMartini, C. 2005. A new and improved design for multiobject iterative auctions. *Management Science* 51(3):419–434.
- Leyton-Brown, K.; Pearson, M.; and Shoham, Y. 2000. Towards a universal test suite for combinatorial auction algorithms. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, 66–76. ACM.
- MacKay, D. J. C. 1992. Bayesian interpolation. *Neural computation* 4(3):415–447.
- Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic theory*. Oxford university Press.
- Minka, T. P. 2001. *A family of algorithms for approximate Bayesian inference*. Ph.D. Dissertation, Massachusetts Institute of Technology.
- Mishra, D., and Parkes, D. C. 2007. Ascending price vickrey auctions for general valuations. *Journal of Economic Theory* 132(1):335–366.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of operations research* 6(1):58–73.
- Nisan, N. 2000. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, 1–12. ACM.
- Opper, M., and Winther, O. 1998. A Bayesian approach to online learning. *Online Learning in Neural Networks* 363–378.
- Parkes, D. C., and Ungar, L. H. 2000. Preventing strategic manipulation in iterative auctions: Proxy agents and price-adjustment. In *Proceedings of the 17th AAAI Conference on Artificial Intelligence*, 82–89.
- Parkes, D. C. 1999. iBundle: an efficient ascending price bundle auction. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, 148–157. ACM.
- Parkes, D. C. 2006. Iterative combinatorial auctions. In Cramton, P.; Shoham, Y.; and Steinberg, R., eds., *Combinatorial auctions*. The MIT Press.
- Petrakis, I.; Ziegler, G.; and Bichler, M. 2012. Ascending combinatorial auctions with allocation constraints: On game theoretical and computational properties of generic pricing rules. *Information Systems Research* 24(3):768–786.
- Sollich, P. 2002. Bayesian methods for support vector machines: Evidence and predictive class probabilities. *Machine learning* 46(1-3):21–52.
- Train, K. E. 2009. *Discrete choice methods with simulation*. Cambridge University Press.
- Williams, C. K., and Rasmussen, C. E. 2006. *Gaussian Processes for Machine Learning*. The MIT Press.