Auction protocols

Vincent Conitzer

The word "auction" generally refers to a mechanism for allocating one or more resources to one or more parties (or **bidders**). Generally, once the allocation is determined, some amount of money changes hands; the precise monetary transfers are determined by the auction process. While in some auction protocols, such as the English auction, bidders repeatedly increase their bids in an attempt to outbid each other, this is not an essential component of an auction. There are many other auction protocols, and we will study some of them in this chapter.

Auctions have traditionally been studied mostly by economists. In recent years, computer scientists have also become interested in auctions, for a variety of reasons. Auctions can be useful for allocating various computing resources across users. In artificial intelligence, they can be used to allocate resources and tasks across multiple artificially intelligent "agents." Auctions are also important in electronic commerce: there are of course several well-known auction websites, but additionally, search engines use auctions to sell advertising space on their results pages. Finally, increased computing power and improved algorithms have made new types of auctions possible—most notably **combinatorial** auctions, in which multiple items are for sale in the same auction, and bidders can bid on bundles of items.

We begin this chapter by studying single-item auctions. Even though most computer scientists are perhaps more interested in combinatorial auctions, single-item auctions allow us to more easily introduce certain concepts that are also of key importance in combinatorial auctions.

1   Standard single-item auction protocols

In this section, we review some basic protocols for auctioning a single item. The reader is encouraged to think about which of these protocols are similar to each other; we will discuss relationships among them shortly.

English. The English auction is the most familiar protocol to most people. In an English auction, every bidder is allowed to place a bid higher than the current highest bid. If at some point, no bidder wishes to place a higher bid, then the bidder with the current highest bid wins the item, and pays her bid.

Dutch. The Dutch auction proceeds in the opposite direction from the English auction. In a Dutch auction, an initial price is set that is very high, after which the price is gradually decreased. At any moment, any bidder can claim the item. She then wins the item and has to pay the current price.

Japanese. In a Japanese auction, the initial price is zero; the price is then gradually increased. A bidder can leave the room when the price becomes too high for her. Once there is only one bidder remaining, that bidder wins the item, and pays the price at which the last other bidder left the room.

First-price sealed-bid. In a first-price sealed-bid auction, each bidder communicates a bid privately to the auctioneer—say, in a sealed envelope. The auctioneer then opens all the envelopes; the bidder with the highest bid wins the item, and pays the bid that she placed.

Second-price sealed-bid (also known as **Vickrey**). The second-price sealed-bid auction proceeds exactly as the first-price sealed-bid auction, except the highest bidder (who still wins the item) now pays the second-highest bid, instead of her own.

Let us consider which of these auction protocols are similar to each other. Perhaps the most obvious similarity is between the English and Japanese auctions. For both, there is a price that is rising, and the last remaining bidder wins. There is a distinction, however: in

an English auction, two bidders may be bidding each other up, while a third bidder quietly sits by, even though she remains interested in the item. In this case, the first two bidders are unaware that they have another competitor. In a Japanese auction, this situation cannot occur.

The Japanese auction and the second-price sealed-bid auction are also closely related (and the English auction is related to the second-price sealed-bid auction in a similar way). Suppose, for a second, that each bidder in a Japanese auction decides at the beginning of the auction on the price at which she will leave the room. Of course, other strategies are possible: a bidder may base how long she stays in the room on which other bidders are still left. However, *if* the bidders follow such strategies, then the bidder who, at the beginning, chose the highest price will end up winning, and she will end up paying the second-highest price selected by a bidder—similarly to the second-price sealed-bid auction.

The Dutch auction and the first-price sealed-bid auction are even more closely related. Similarly to the Japanese auction, in a Dutch auction, a bidder may decide at the beginning on the price at which she will claim the item (if this price is reached). In fact, unlike in the Japanese auction, there is little else that a bidder can do in terms of strategizing. In a Japanese auction, a bidder can let her bidding strategy depend on who else is left; but in a Dutch auction, there is nothing to condition her strategy on, since the only event that can happen is that someone else claims the item—but at that point the auction is over and it no longer matters what anyone does. Now, the bidder that chooses the highest price will win, and pay that price—similarly to the first-price sealed-bid auction. Because of this argument, the Dutch and first-price sealed-bid auctions are usually considered strategically equivalent.

We will see some other single-item auctions later in this chapter. For most of this chapter, we will focus on sealed-bid auctions. As we have seen, for each of the auctions studied so far, there is a roughly equivalent sealed-bid auction; in a sense, this is true for *any* auction, as we will see shortly. Nevertheless, there are reasons to use English, Japanese,

and Dutch auctions (more generally, **ascending** and **descending** auctions). One reason is that they allow bidders to postpone certain decisions until later. For example, if a bidder in a Dutch auction is deciding whether to claim the item at $60 or $50, she may as well wait until the price drops to, say, $70, before she starts to think about what she will do. If another bidder claims the item before that, the former bidder will have saved herself some unnecessary agonizing. This is related to **preference elicitation**, which we will discuss towards the end of this chapter.

2   Valuations and utilities

How a bidder should bid in an auction depends on how much the item for sale is worth to her. In this chapter, we will assume that each bidder can determine how much the item is worth to her, and that events in the auction will not change her assessment. That is, each bidder i has an unchanging **valuation** v_i for the item, which she can determine at the beginning of the auction. This assumption is not always realistic. For example, if a bidder sees that other bidders are bidding aggressively on an item, this may be evidence to her that, before the auction, those bidders inspected the item in person and found it to be of good quality. This evidence may improve the first bidder's perception of—and hence, valuation for—the item. Settings such as these, where some bidders have private information that would affect the valuation of other bidders for the item, are known as **interdependent valuations** settings. Most research assumes away the possibility of interdependent valuations, and we will do so in this chapter.

In general, each bidder has a **utility** for each outcome of the auction, and acts to maximize her expected utility. We will assume that a bidder's utility for winning the item and having to pay π_i is $u_i = v_i - \pi_i$, and her utility for not winning the item and having to pay π_i is $u_i = -\pi_i$. (Generally, losing bidders will not be made to pay anything; however, later in this chapter, we will see an auction that makes payments *to* losing bidders, in which

case π_i is negative.) Thus, we are assuming that a bidder's utility function decomposes into separate valuation and payment components, and that utility is linear in money. This assumption is known as the **quasilinear preferences** assumption. It, too, is not always realistic, for the following reasons. In general, one's utility may be strictly concave in money (that is, one may have **decreasing marginal utility** for money: the utility of having another (say) dollar may decrease as one accumulates more money), since at some point one runs out of uses for money. Also, in general, the effect of money on utility may depend on whether one has won the item: for example, if a bidder wins a pair of skis in an auction, she needs money to travel on a skiing vacation (and hence has high marginal utility for money), whereas if she does not win, she has less use for additional money (and hence has low marginal utility for money). Nevertheless, the quasilinearity assumption is usually made, and we will do so in this chapter.

Another assumption that is implicit in the above is that a bidder that does not win the item does not care about which other bidder wins the item, and that bidders do not care about how much other bidders pay. This assumption is known as the **no externalities** assumption. Once again, this assumption is not always realistic: a bidder may prefer to see the item end up with a friend rather than with an enemy, or she may prefer to see the other bidders run out of money so that they will not compete in future auctions. Again, we will not go into detail on this in this chapter.

3 Strategic bidding

It does not always make sense for a bidder to simply bid her true valuation. For example, if bidder i bids her true valuation v_i in a first-price sealed-bid auction, then even if she wins, her utility will be $v_i - \pi_i = v_i - v_i = 0$. Hence, she should bid lower than her true valuation to have any chance of obtaining positive utility. But how much lower should she bid? Intuitively, this should depend on her beliefs about the other bidders' valuations. If

she expects to be the only bidder who is seriously interested in obtaining the item, she can place a low bid and still probably win; whereas if she expects there to be many other competitive bidders, she should bid closer to her true valuation to have a decent chance of winning. However, it is not obvious how to calculate her optimal bid precisely, even given a probability distribution over the others' valuations. This is because she cannot expect the other bidders to bid their true valuations, either: they also need to bid below their true valuations to have a chance of obtaining positive utility. And precisely how much lower is optimal for them to bid depends, in turn, on how the first bidder bids.

3.1   Solving the first-price sealed-bid auction

To resolve this circularity, we need to turn to **game theory**, which studies settings in which each bidder's (or, more generally, **agent's**) optimal course of action depends on the actions of the other bidders. To apply game theory to the first-price sealed-bid auction, we first need to introduce the concept of a **strategy**. In a sealed-bid auction, a strategy for bidder i is a function $s_i : \mathbb{R}^{\geq 0} \to \mathbb{R}^{\geq 0}$, where $s_i(v_i)$ is the bid that i will place if her true valuation is v_i. That is, for *every* valuation that the bidder may have, the strategy specifies what she should bid. This may appear somewhat excessive: if the bidder already knows her true valuation, why should she have to specify what she would have bid if her valuation had been different? The reason that we need to think about this is that the *other* bidders do not know bidder i's valuation, and *they* need to think about what i would do for each valuation. In turn, bidder i needs to think about what the other bidders will do, and hence also needs to think about what they think she will do.

Let us suppose that for each bidder i, there is a (commonly known) prior probability distribution p_i over her valuation v_i; moreover, let us assume that the valuations are drawn independently. Hence, each bidder i knows her own valuation v_i exactly, but for every other bidder $j \neq i$, i's probability distribution over j's valuation is p_j. Now, *if* bidder i knows the strategies of the other bidders, then for every bid that she might place, she can evaluate her

6

expected utility; and of course she should choose one that maximizes her expected utility. As is typically done in game theory, we will look for an **equilibrium**, which prescribes a strategy for every bidder such that, for every bidder, for every possible valuation for that bidder, her strategy will prescribe a bid that maximizes her expected utility, given the other strategies. Formally, a **Bayes-Nash equilibrium** consists of a strategy $s_i : \mathbb{R}^{\geq 0} \to \mathbb{R}^{\geq 0}$ for every bidder such that, for every bidder i, for every $v_i \in \mathbb{R}^{\geq 0}$, and for every alternative bid $\hat{v}_i \in \mathbb{R}^{\geq 0}$,

$$\int_{v_{-i}} (\prod_{j \neq i} p_j(v_j)) u_i(v_i, s_{-i}(v_{-i}), s_i(v_i)) dv_{-i} \geq \int_{v_{-i}} (\prod_{j \neq i} p_j(v_j)) u_i(v_i, s_{-i}(v_{-i}), \hat{v}_i) dv_{-i}$$

Let us dissect this complicated inequality. First, the notation $-i$ is shorthand for "the bidders other than i," so that v_{-i} is shorthand for $v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n$. The notation $\hat{v}_i$ is generally used for bidder i's bid, not necessarily equal to her true valuation v_i. $u_i(v_i, \hat{v}_{-i}, \hat{v}_i)$ is the utility that bidder i obtains if her true valuation is v_i, but she bid $\hat{v}_i$, and the other bidders bid $\hat{v}_{-i}$. In the first-price sealed-bid auction, $u_i(v_i, \hat{v}_{-i}, \hat{v}_i) = v_i - \hat{v}_i$ if $\hat{v}_i$ is higher than all the bids in $\hat{v}_{-i}$, and it is 0 otherwise. Now we can see that the inequality says that, if the bidders other than i follow their strategies, then i's expected utility for bidding $s_i(v_i)$ should be at least equal to her expected utility for bidding any other $\hat{v}_i$—that is, she should not be able to do better by not following the strategy s_i, given that the other bidders are indeed following their strategies s_{-i}.

We will now give an example of an equilibrium. Suppose that each p_i is a uniform distribution over $[0, 1]$. We will show that the strategies defined by $s_i(v_i) = v_i(n-1)/n$ (where n is the number of bidders) constitute an equilibrium. (We will not get into detail here on how one might have actually *derived* these strategies, but there are techniques for doing so.) Suppose that all other bidders $(-i)$ indeed follow these strategies. Then, the expected utility for i of bidding $\hat{v}_i \leq (n-1)/n$ is $(\hat{v}_i n/(n-1))^{n-1}(v_i - \hat{v}_i)$, because the probability that a given other bidder bids less than $\hat{v}_i$ is $\hat{v}_i n/(n-1)$. (There is no reason

to bid more than $(n-1)/n$, because no other bidder will bid more than $(n-1)/n$.) Using simple calculus, one can check that this expression is maximized by setting $\hat{v}_i$ equal to $v_i(n-1)/n$—exactly as the strategy prescribes! This proves that these strategies indeed constitute an equilibrium.

3.2   Solving the second-price sealed-bid auction

Now, let us turn to the second-price sealed-bid auction. As it turns out, the analysis needed to solve this auction is not nearly as complicated. In fact, in the second-price sealed-bid auction, it is always optimal for a bidder to bid her true valuation, regardless of the other bids! That is, the strategy $s_i(v_i) = v_i$ is a **dominant strategy**. While this may come as a surprise at first, it is not so difficult to see why it is true. Suppose, for a second, that bidder i can actually see the others' bids before placing her own bid. Let us consider the value $\hat{v}_{max}$, the highest bid among the other bidders. Bidder i effectively has only two choices: to bid higher than $\hat{v}_{max}$, and obtain utility $v_i - \hat{v}_{max}$; or to bid lower than (or the same as) $\hat{v}_{max}$, and obtain utility 0. Clearly, she should do the former if and only if $v_i > \hat{v}_{max}$. But this is exactly what would happen if she just bid her true valuation—for which she does not even need to know the others' bids! Hence, by bidding her true valuation, she performs as well as she could have performed even if she had known the others' bids. This implies that bidding truthfully is also a Bayes-Nash equilibrium, although the result is much stronger than that. For example, in the first-price auction, if one knows the bids of the other bidders, then certainly one might be better off bidding differently from the equilibrium that we derived for that auction (which we derived under the assumption that bidders do not know each other's valuations). Hence, the strategies in the first-price auction equilibrium are not dominant strategies. Mechanisms in which revealing one's true valuation is a dominant strategy (such as the second-price sealed-bid auction) are called **(dominant-strategies) incentive compatible**, **strategy-proof**, or simply **truthful**.

4 Revenue equivalence

Now that we have analyzed how bidders should bid in these two auctions, let us ask the following question: which one obtains more revenue for the seller, in expectation? The answer is not immediately obvious: naïvely, one might say that the first-price auction should result in more revenue, since after all it charges the highest bid rather than the second-highest; but then again, in equilibrium, the bids are lower in the first-price auction. Which of these two effects is stronger?

For the case of independent uniform priors over $[0, 1]$, we can compute the expected revenues using the equilibrium strategies from above. For the first-price auction, the probability that all bids are below a given value b is $bn/(n - 1)$, which is also the probability that the revenue will be below b. That is, this expression gives the cumulative density function of the revenue of the first-price auction, and using it one can compute the expected revenue to be $(n - 1)/(n + 1)$. For the second-price auction, the probability that there is at most one bid higher than b is $b^n + nb^{n-1}(1 - b)$, which is also the probability that the revenue will be below b. That is, this expression gives the cumulative density function of the revenue of the second-price auction, and using it one can compute the expected revenue to be $(n - 1)/(n + 1)$—the same as that for the first-price auction! This is no accident: it is a special case of the following result, which is known as the **revenue equivalence theorem** [Myerson, 1981; Riley and Samuelson, 1981].

Theorem 1 *Suppose that the bidders' valuations are independent and identically distributed over a continuous interval $[L, H]$, and that there are no "gaps" in this distribution. Then, any two auction mechanisms that*

1. *in equilibrium always allocate the item to the bidder with the highest valuation, and*

2. *give a bidder with valuation L an expected utility of 0,*

will result in the same expected revenue for the seller.

(There are more general versions of this result.) In the next section, we will see single-item auction mechanisms that result in different expected revenues, because they violate one of the two conditions in the theorem. From this point on, we will study only truthful mechanisms. This is justified by a result known as the **revelation principle** [Gibbard, 1973; Green and Laffont, 1977; Myerson, 1979, 1981], which states (roughly) that, if bidders bid strategically, then for every mechanism that is not a truthful mechanism, there is a truthful mechanism that performs equally well.

5 Auctions with different revenues

Suppose that we have a prior distribution over each bidder's valuation, and that we wish to design a mechanism that maximizes expected revenue. It is easy to see that running, say, a second-price sealed-bid auction is not always optimal. For example, suppose there is only one bidder. The second-price sealed-bid auction will never collect any revenue in this case, because there is no second bidder. However, we can also make a **take-it-or-leave-it offer** to the one bidder: the bidder will obtain the item if and only if it is worth more than some fixed value k to her, in which case she will pay k; otherwise, the seller keeps the item. This will generate a revenue of k at least some of the time. While this may not seem like an auction, setting a **reserve price** of k in a second-price sealed-bid auction will have the same effect. (In such an auction, if only one bid is above the reserve price, then that bid pays the reserve price.)

In general, the auction that maximizes expected revenue is known as the **Myerson auction** [Myerson, 1981], and it proceeds as follows. For each bidder i, compute her **virtual valuation** $\psi_i(\hat{v}_i)$ as a function of her bid, as follows:

$$\psi_i(\hat{v}_i) = \hat{v}_i - (1 - F_i(\hat{v}_i))/f_i(\hat{v}_i)$$

Here, F_i is the cumulative density function of i's valuation, and f_i is its derivative, the

probability density function. The bidder with the highest virtual valuation wins, unless this bidder has a virtual valuation below 0, in which case nobody wins. The price that the winning bidder pays is the lowest value that she could have bid while still winning. For example, if each bidder's valuation is drawn from the uniform distribution over $[0, 1]$, then the Myerson auction becomes a second-price sealed-bid auction with a reserve price of $1/2$. This is because

$$\psi_i(1/2) = 1/2 - (1 - F_i(1/2))/f_i(1/2) = 1/2 - (1 - 1/2)/1 = 0$$

It does not always make sense to try to maximize expected revenue. In some settings, our main goal is to allocate the item efficiently (that is, to the bidder that values it most), and payments are merely a necessary nuisance in achieving this goal. For example, suppose that several parties jointly own an item, and they wish to run an auction amongst themselves to decide on a single owner. What should happen to the revenue of this auction? It seems to make sense to redistribute it back to the bidders themselves, but doing so effectively changes the auction mechanism. For example, suppose that the bidders run a second-price sealed-bid auction for the item, and then redistribute the revenue of this auction equally (each bidder receives $1/n$ of the revenue). Unlike the second-price sealed-bid auction without redistribution, this auction is actually *not* truthful: the second-highest bidder now has an incentive to increase her bid to drive up the price that the highest bidder pays, because the second-highest bidder will receive a fraction of this price.

Fortunately, it turns out that we can redistribute at least some of the revenue while maintaining truthfulness. One auction mechanism that achieves this is the following (independently invented on at least three different occasions [Bailey, 1997; Porter *et al.*, 2004; Cavallo, 2006]). Let us define v_{-i}^2 to be the second-highest bid among bidders *other than* bidder i. For the top two bidders, this is the third-highest bid overall (v^3); for the remaining $n - 2$ bidders, it is the second-highest bid (v^2). We run the second-price sealed-bid

11

auction, and we redistribute v_{-i}^2/n to bidder i. This redistribution payment does not affect bidders' incentives in bidding, because no bidder can affect her own redistribution payment. Hence, the auction remains truthful. Additionally, the total redistributed is $2v^3/n + (n-2)v^2/n \leq v^2$—so the total redistributed is no more than is collected from the second-price sealed-bid auction. A total of $2(v^2 - v^3)/n$ is not redistributed. This money must be given to someone else (but not someone that the bidders care about, since that might affect their incentives), or, say, burned. It is impossible to achieve efficient allocation without ever wasting any money, but it is possible to waste even less money (either on average or in the worst case): this is achieved by also letting bidder i's redistribution payment depend on $v_{-i}^3, v_{-i}^4, \ldots, v_{-i}^{n-1}$ [Guo and Conitzer, 2007; Moulin, 2007].

It is interesting to note that the revenue equivalence theorem from above does not apply to Myerson's auction because the first condition is not satisfied; it does not apply to the redistribution mechanisms because the second condition is not satisfied.

6 Complementarity and substitutability

Now that we have studied single-item auctions, let us consider settings where multiple items are for sale. One possibility is to sell each item in a separate single-item auction; these auctions can be held simultaneously (**parallel** auctions) or back-to-back (**sequential** auctions). *If*, for each item, each bidder's valuation for that item does not depend on which other items she wins, then the individual auctions are entirely separate events, and we can apply the techniques that we have studied up to this point. However, this is not always a realistic assumption. For example, if the items for sale are a plane ticket to Las Vegas and a hotel reservation in Las Vegas, then it may be that the bidder's valuation for the plane ticket alone is 200, her valuation for the hotel reservation alone is 100, but her valuation for both together is 500. The package of both items is worth more than the sum of its parts, that is, the items are **complementary**.

Now suppose that the plane ticket is auctioned first, and the hotel reservation second (both in second-price sealed-bid auctions). How much should the bidder bid in the first auction? If she bids 200 for the ticket, she may lose to a bidder bidding 201, only to later find out she could have won the hotel reservation for 101, so that she regrets not bidding higher in the first auction. However, if she bids 400 for the ticket, she may win it at a price of 399, only to later find out that the hotel reservation sells for 1000, so that she regrets winning the ticket. It is not clear what the bidder should do—she no longer has a dominant strategy. Moreover, the resulting allocation of items can be inefficient.

Another possibility is that the package of items is worth *less* than the sum of its parts, in which case the items are said to be **substitutable**. For example, if reservations for two different hotels are for sale, a bidder may value each individual reservation at 100, but the package for both reservations at 150. In sequential auctions, substitutability can cause problems similar to those caused by complementarity. Both can also cause similar problems in parallel auctions.

Instead of making the bidders agonize over the prices at which items in future or parallel auctions are likely to sell, an alternative is to let each bidder report *all* her valuations, one for each subset of the items, and decide on the allocation of items to bidders based on that information. This is what is done in a combinatorial auction, and it circumvents the problems that parallel and sequential auctions run into when there are complementarities and substitutabilities.

7 Combinatorial auctions

In a **combinatorial auction**, a set I of multiple items is (simultaneously) for sale, and bidders can bid on any **bundle** (that is, subset) of items. If we again make the assumptions that bidders' valuations for the items do not change based on other bidders' private information, utilities are quasilinear, and there are no externalities, then each bidder i has

a privately held valuation function $v_i : 2^I \to \mathbb{R}^{\geq 0}$, where $v_i(S)$ is i's valuation for bundle $S \subseteq I$; and the utility of bidder i when she wins bundle S and pays π_i is $v_i(S) - \pi_i$. Generally it is assumed that $v_i(\emptyset) = 0$, and additionally that for $S \subseteq S'$, $v_i(S) \leq v_i(S')$. (The latter assumption is often called **free disposal**: receiving additional items can never decrease a bidder's valuation, because at worst the additional items can simply be discarded.) We will start by looking at sealed-bid combinatorial auctions. An immediate problem with this approach is that in general, each bidder must reveal $2^m - 1$ real numbers (where $m = |I|$), one for each nonempty bundle. Once m gets to be somewhat large, this becomes impractical. However, there is usually some structure in the bidders' valuation functions, so that they can be represented more concisely.

One very restrictive, but commonly studied assumption about this structure is that bidders are **single-minded**. A bidder i is single-minded if there exists some bundle S_i and some real number v_i such that $v_i(S) = v_i$ if $S_i \subseteq S$, and $v_i(S) = 0$ otherwise. That is, there is a single bundle of items that the bidder wishes to obtain; she will simply discard any additional items, and if she fails to obtain even one item within her desired bundle, her valuation drops to 0. If bidders are single-minded, then a bid can be represented simply as an ordered pair $(S_i, v_i) \in 2^I \times \mathbb{R}^{\geq 0}$.

A single-minded bid cannot represent even fairly straightforward valuation functions, such as **additive** valuation functions. (A valuation function v_i is additive if for all $S \subseteq I$, $v_i(S) = \sum_{s \in S} v_i(\{s\})$. That is, there are no complementarities or substitutabilities.) We would like to give bidders some more flexibility, by providing them with a richer **bidding language** in which to describe their valuation function. One such bidding language is the **OR language**, which effectively allows a bidder to submit multiple single-minded bids. Formally, a bid in the OR language takes the form (S_1, v_1) OR (S_2, v_2) OR ... OR (S_k, v_k). Such a bid is interpreted as follows: for any subset $T \subseteq \{1, \ldots, k\}$ with the property that for any $j_1, j_2 \in T, j_1 \neq j_2$, we have $S_{j_1} \cap S_{j_2} = \emptyset$, $v_i(\bigcup_{j \in T} S_j) = \sum_{j \in T} v_j$. That is, the auctioneer can accept any subcollection of the single-minded bids within the OR-bid, as

long as there is no overlap between the accepted S_j. (To be precise, the last $=$ symbol should really be a $\geq$ symbol, because it is possible that the same set $\bigcup_{j \in T} S_j$ (or a subset thereof) can be written as a union of disjoint S_j in a different way that results in a greater sum of v_j. The bidder's valuation for the subset is the maximum value that can be obtained in this way. Hence, equality holds only if there is no better way to write the bundle as a union of disjoint S_j. For example, given the bid $(\{A\}, 1)$ OR $(\{A, B\}, 3)$ OR $(\{B, C\}, 3)$ OR $(\{C\}, 2)$, we have $v_i(\{A, B, C\}) = v_i(\{A, B\}) + v_i(\{C\}) = 5 > 4 = v_i(\{A\}) + v_i(\{B, C\})$.) The OR language allows for representing additive valuations, simply by OR-ing together singleton bundles.

Nevertheless, there are valuation functions that the OR language cannot capture. For example, suppose there are two items, a and b, and that bidder i's valuation function is given by $v_i(\{a\}) = 1, v_i(\{b\}) = 1, v_i(\{a, b\}) = 1$. That is, she wants either item, and having both items is no more useful to her than having a single one. This function cannot be represented in the OR language: the bid would have to contain the terms $(\{a\}, 1)$ and $(\{b\}, 1)$, but this would already imply that $v_i(\{a, b\}) \geq 2$. A language that *can* capture this valuation function is the **XOR language**. The difference between the OR and XOR languages is that the auctioneer can accept at most *one* of the single-minded bids that are XORed together, even if they do not overlap. For example, the above valuation function is easily expressed as $(\{a\}, 1)$ XOR $(\{b\}, 1)$: this bid implies a valuation of only 1 for the bundle $\{a, b\}$, since it is not possible to accept both single-minded bids in the bid. Using XORs, we can in fact represent *any* valuation function, by using a single-minded bid for every possible bundle and XOR-ing them all together. Of course, this is not a very concise representation. Unfortunately, even representing additive valuation functions can require exponentially long bids if we use only XORs. But it is also possible to use ORs and XORs simultaneously, to get the best of both worlds. For example, the bid $((\{a\}, 1)$ XOR $(\{b\}, 1))$ OR $(\{c\}, 2)$ indicates a value of $1 + 2 = 3$ for the bundle $\{a, b, c\}$. There are other bidding languages that are not based on ORs and XORs, but we will not discuss them in this

chapter.

8 The winner determination problem

Now that we have considered how to bid in a (sealed-bid) combinatorial auction, we must consider how to determine who wins what. Of course, we cannot award a single item to two different bidders. But this still leaves plenty of options. One natural approach is to maximize **efficiency**, the total value generated. That is, if S_i is the bundle of items that we award to bidder i, then we should maximize $\sum_{i=1}^{n} \hat{v}_i(S_i)$ (under the constraint that $S_i \cap S_j = \emptyset$ for all $i \neq j$). This optimization problem is known as the **winner determination problem (WDP)**.

As it turns out, the WDP is computationally hard. Even if we consider the special case where each bidder submits a single-minded bid, the problem turns out to be equivalent to WEIGHTED-SET-PACKING, which is NP-hard [Rothkopf *et al.*, 1998] and inapproximable [Sandholm, 2002]. On the other hand, the problem does not become any harder if we allow bidders to use ORs, since a bidder using ORs is effectively submitting multiple single-minded bids. In fact, even if we allow XORs, the problem in a sense becomes no harder: this is because, for the purpose of solving the WDP, we can transform an instance with XORs into one with only ORs using the following trick [Fujishima *et al.*, 1999; Nisan, 2000]. Given a bid of the form (S_1, v_1) XOR (S_2, v_2), we create a new "dummy" item, d, and replace the bid by $(S_1 \cup \{d\}, v_1)$ OR $(S_2 \cup \{d\}, v_2)$. Even though there is an OR between these two bids, they cannot both be accepted, since they have an item in common; moreover, since no other bids mention this item, everything else remains unaffected. Because of this, in the next two subsections, we will focus on algorithms for the single-minded case only. We will consider optimal algorithms for both the general case and special cases. We will postpone discussion of approximation algorithms until later, because we will require a particular kind of approximation algorithm in this context.

8.1 General-purpose winner determination algorithms

Given single-minded bids $\{(S_i, \hat{v}_i)\}$, one straightforward way to solve the winner determination problem is to solve the following integer program, which uses a binary variable b_i to indicate whether bid i is accepted:

maximize $\sum_i b_i \hat{v}_i$

subject to

for each $s \in I, (\sum_{i:s \in S_i} b_i) \leq 1$

for each $i, b_i \in \{0, 1\}$

The main constraint of this integer program ensures that each item is awarded at most once. Software packages such as CPLEX can be used to solve such an integer program; since the WDP is NP-hard, it should come as no surprise that solving integer programs is also NP-hard.

One interesting aside is that in some settings, bids can be accepted **partially**. For example, if we have three bids, $(\{a, b\}, 2)$, $(\{a, c\}, 2)$, and $(\{b, c\}, 2)$, it may be possible to accept *half* of each bid, awarding half of a and half of b to the first bidder for a value of 1, half of a and half of c to the second bidder for a value of 1, and half of b and half of c to the third bidder for a value of 1. This gives us a total value of 3 (we note that if we cannot accept bids partially, then we can obtain only 2). If accepting bids partially is possible, then we can solve the WDP by modifying the above program slightly. We make the b_i continuous variables, replacing the last constraint by

for each $i, 0 \leq b_i \leq 1$

At this point, the program has become a linear program, and linear programs can be solved in polynomial time [Khachiyan, 1979]. However, in the remainder of this chapter we will assume that it is not possible to accept bids partially.

An alternative approach to solving the general WDP is to write a search algorithm

based on techniques from artificial intelligence; for an overview of work along this line, see [Sandholm, 2006]. It should be noted that such algorithms are in many ways similar to algorithms for solving integer programs.

Yet another option is to use the following dynamic programming approach [Rothkopf *et al.*, 1998]. For any subset $S \subseteq I$, let $w(S)$ be the maximum total value that can be obtained using only items in S (that is, if we threw away the items in $I - S$). Let $B(S)$ be the collection of all proper subsets $S' \subset S$ such that there is at least one bid on exactly S'. Then, we have:

$$w(S) = \max\{\max_i \hat{v}_i(S), \max_{S' \in B(S)} w(S') + w(S - S')\}$$

Since the occurrences of w on the right-hand side involve subsets smaller than S, we can use dynamic programming to compute $w(S)$ for every subset, starting with the small ones and working our way up to I—and $w(I)$ gives the maximum value that can be obtained overall. This algorithm runs in $O(n3^m)$ time (where m is the number of items). It is straightforward to extend the dynamic program to keep track not only of the values that can be obtained, but also of the bids that need to be accepted to obtain these values.

8.2 Special-purpose winner determination algorithms

The WDP is NP-hard in general. Nevertheless, if the bids have some structure, then the WDP is sometimes solvable in polynomial time. For example, suppose that there are only (single-minded) bids on *pairs* of items [Rothkopf *et al.*, 1998]. Certainly, if multiple bids bid on the same pair of items, it never makes sense to accept a bid that is not the highest. Hence, we know the value of pairing any two given items together for sale (namely, the highest bid for that pair), and the only decision left is which items to pair together. This is a MAXIMUM-WEIGHTED-MATCHING problem, which can be solved in polynomial time. (This can easily be extended to also allow for bids on individual items—for example, by adding dummy items.) Unfortunately, if we allow for bids on sets of three items, the

problem becomes NP-hard (by reduction from EXACT-COVER-BY-3-SETS).

As another example, suppose that the items are arranged as the vertices of a graph, and that every bid is on a bundle of items that constitutes a *connected component* in the graph. This is always possible by adding an edge between every pair of items (that is, making the graph a complete graph), but we will be interested in restricted classes of graphs. In particular, if the graph is a tree or a cycle, or, more generally, has **bounded treewidth**, then the WDP can be solved in polynomial time using dynamic programming [Sandholm and Suri, 2003; Conitzer *et al.*, 2004]. To use a result like this, we can either collect the bids first and then find a graph with which they are consistent, or we can specify the graph beforehand and require all bids to be consistent with this graph. A later result generalizes this even further by considering **hypertree decompositions** [Gottlob and Greco, 2007].

There is various other work on structure that bids may have that makes the WDP easier [Tennenholtz, 2000; Penn and Tennenholtz, 2000; Sandholm and Suri, 2003]. Even if the bidders' valuations are not likely to have the required structure exactly, one possibility is to *force* the bidders to only use bids with this structure. This comes at the loss of some economic efficiency, because bidders can no longer express their exact valuations; nevertheless, it is generally better than reverting to single-item auctions.

9   The Generalized Vickrey Auction

So far, we have not yet considered how much a winning bidder should pay in a combinatorial auction. We could simply make such a bidder pay her bid (that is, her reported valuation for the bundle she won), resulting in a first-price sealed-bid combinatorial auction. As in the case of a first-price sealed-bid single-item auction, no bidder would ever bid her true valuation function, since this would guarantee that even if she wins something, she will have a utility of 0. We recall that an auction is truthful if revealing one's true valuation function is always optimal, regardless of the others' bids. Can we create a truthful combinatorial

auction? A natural approach is to try to generalize the (truthful) second-price sealed-bid (aka. Vickrey) auction to the combinatorial setting. There are multiple ways in which one can generalize the Vickrey auction. For example, one can charge a winning bidder the highest other bid that was placed on the exact same bundle. To see that this is not a good idea, consider the following example. Suppose that the bids are $(\{a\}, 1)$, $(\{b\}, 1)$, and $(\{a, b\}, 5)$, all from different bidders. The third bidder would win both items, and this bidder would pay 0, because nobody else bid on the bundle $\{a, b\}$. This intuitively feels wrong, since there was demand for the items from the other bidders. If the third bidder had bid $(\{a\}, 5)$ instead, she would have had to pay 1; thus, by bidding for *more* items, the bidder actually pays *less*! This also shows that this particular generalization is not truthful: if the third bidder is in fact interested only in a, she is still better off bidding on both items, and just throwing b away.

Fortunately, there is another generalization that does work. Let $\hat{V}$ be the total value of the accepted bids. Let $\hat{V}_{-i}$ be the total value that would have resulted if i had never entered the auction. (Computing this requires solving the winner determination problem again, this time without i.) Then, if S_i is the bundle that bidder i wins (possibly the empty bundle), she must pay $\hat{V}_{-i} - (\hat{V} - \hat{v}_i(S_i))$. This expression is the difference between how much the other bidders would have valued the allocation that would have resulted if i had never been present, and how much they value the current allocation. (For this reason, it is sometimes said that i pays the **externality** that she imposes on the other bidders.) Let us consider the above example with bids $(\{a\}, 1)$, $(\{b\}, 1)$, and $(\{a, b\}, 5)$. If the third bidder had not been present, the first two bids would have been accepted, so $V_{-3} = 2$. Hence, bidder 3 pays $2 - (5 - 5) = 2$. Let us make the example slightly richer, by adding bids $(\{c\}, 2)$ and $(\{a, c\}, 5)$, again from different bidders. Now, the third and fourth bidders win, for a total value of $5 + 2 = 7$. Without the third bidder, the second and fifth bidders would have won, for a total value of $1 + 5 = 6$. Hence, the third bidder must pay $6 - (7 - 5) = 4$. Without the fourth bidder, again, the second and fifth bidders would have won, for a total

value of 6. Hence, the fourth bidder must pay $6 - (7 - 2) = 1$.

This way of computing payments is usually called the **Generalized Vickrey Auction (GVA)**. It is sometimes also called the **Clarke** mechanism, or the **VCG** mechanism (for Vickrey, Clarke, and Groves [Vickrey, 1961; Clarke, 1971; Groves, 1973]). (Clarke and VCG refer to generalizations beyond auctions.) The GVA has several nice properties. For one, it is truthful. To see this, we note that bidder i's eventual utility is $v_i(S_i) - \pi_i = v_i(S_i) - (\hat{V}_{-i} - (\hat{V} - \hat{v}_i(S_i))) = v_i(S_i) + (\hat{V} - \hat{v}_i(S_i)) - \hat{V}_{-i}$. It is impossible for i to affect $\hat{V}_{-i}$; therefore, i can focus on maximizing $v_i(S_i) + (\hat{V} - \hat{v}_i(S_i)) = v_i(S_i) + \sum_{j \neq i} \hat{v}_j(S_j)$. There is only one way in which this expression depends on i's bid $\hat{v}_i$: her bid affects the chosen allocation $S_1, \ldots, S_n$. Now, if i had complete control over the chosen allocation (which she does not, but let us suppose for a second that she does), then she would choose the S_j to maximize $v_i(S_i) + \sum_{j \neq i} \hat{v}_j(S_j)$. The winner determination algorithm, on the other hand, chooses the S_j to maximize $\sum_j \hat{v}_j(S_j) = \hat{v}_i(S_i) + \sum_{j \neq i} \hat{v}_j(S_j)$. The only difference between the two expressions is that the first uses v_i, and the second uses $\hat{v}_i$. But then, if bidder i truthfully reports $\hat{v}_i = v_i$, the two expressions will be the same, and the winner determination algorithm will choose exactly the allocation that maximizes i's utility! Hence, the GVA is truthful.

A very observant reader may have noticed that in the proof of truthfulness, the only property of the payment term $\hat{V}_{-i}$ that we used is that i cannot affect it with her bid. Therefore, if truthfulness is all that we care about, we can replace the term $\hat{V}_{-i}$ in the payment expression with any other term that does not depend on i's bid. The mechanisms that can be obtained in this way are the **Groves** mechanisms (the G in VCG). The GVA mechanism, however, does have some additional nice properties that not all Groves mechanisms have. For one, it satisfies **voluntary participation** (aka. **individual rationality**): a bidder never receives negative utility as a result of participating in the auction, as long as she bids truthfully. This is because $\hat{V} \geq \hat{V}_{-i}$ (if this were false, it would mean that we had chosen a suboptimal allocation), and hence $v_i(S_i) - \pi_i = v_i(S_i) + (\hat{V} - \hat{v}_i(S_i)) - \hat{V}_{-i} \geq v_i(S_i) - \hat{v}_i(S_i)$;

and the last expression is zero if i reports truthfully. A final nice property of the GVA is the **nondeficit** (aka. **weak budget balance**) property: at least as much money is collected from the bidders as is given to them. In fact, no money is ever given to a bidder. This is because $\hat{V}_{-i} \geq \hat{V} - \hat{v}_i(S_i)$ ($\hat{V} - \hat{v}_i(S_i)$ is the value of an allocation that is feasible even when i is not present), and hence $\pi_i = \hat{V}_{-i} - (\hat{V} - \hat{v}_i(S_i)) \geq 0$.

10 Collusion and false-name bidding

The GVA is truthful, so it is not possible for an individual bidder to benefit from misreporting her valuation function. However, if multiple bidders simultaneously misreport, that is, they **collude**, then it is possible that all of them benefit from this. To some extent, this problem occurs even in a single-item Vickrey auction. For example, suppose that there are three bidders with valuations 1, 3, and 4. If the third bidder can convince the second bidder not to place any bid, then the third bidder has to pay only 1 instead of 3. The third bidder may even pay the second bidder 1 for staying out, so that they each increase their utility by 1 from this. (In general, the colluders need some protocol for colluding [Graham and Marshall, 1987; Leyton-Brown *et al.*, 2000, 2002], but we will not concern ourselves with that here.)

Still, there is a limit on what colluders can achieve in a single-item Vickrey auction. For example, they can never win the item at a price lower than the highest bid by a noncolluder; nor can they reduce the seller's revenue below what the seller would have made if the colluders had not participated. It turns out that in a combinatorial auction, not even these properties are true. For example, consider a GVA with only two items, a and b. Suppose two bidders have each placed a bid $(\{a, b\}, 1)$. If these are the only bidders, then one of them will win, and pay 1. However, let us now suppose that there two additional bidders (the colluders): one of them bids $(\{a\}, 2)$, and the other bids $(\{b\}, 2)$. The colluders then win. How much does each colluder pay? If we remove one of the colluders, then the

other colluder still wins—that is, the remainder of the allocation does not change. Thus, each colluder (individually) imposes no externality on the other bidders, and hence pays 0. The colluders benefit from this (assuming they have some value for the items they receive), and the auction's revenue has actually decreased as a result of the additional bids. More detail on collusion in the GVA can be found, for example, in [Ausubel and Milgrom, 2006] and [Conitzer and Sandholm, 2006].

The same example can also be used to demonstrate a different vulnerability of the GVA. Suppose that the auction is run in an open, anonymous environment such as the Internet. In such an auction, it is usually possible for a single bidder to participate in the auction under multiple identifiers (for example, e-mail addresses). Thus, given two other bids $(\{a, b\}, 1)$, a single bidder can place a bid $(\{a\}, 2)$ under one identifier, and a bid $(\{b\}, 2)$ under another identifier. As a result, the "false-name" bidder will win both items, and, as before, the price charged to each bid is 0. Hence, bidders sometimes have an incentive to bid under multiple identifiers; that is, the GVA is not **false-name-proof** [Yokoo *et al.*, 2001, 2004].

It should be emphasized that this manipulation cannot be performed simply by submitting multiple bids under a single identifier (for example by bidding $(\{a\}, 2)$ OR $(\{b\}, 2)$). In this case, to compute the bidder's GVA payment, we would remove the bidder's OR-bid in its entirety, so that the resulting payment would be 1. In the example where the bidder uses two different identifiers, there would be no problem if we could tell that the two identifiers correspond to the same real bidder, because in that case we would remove both identifiers' bids simultaneously to compute the bidder's GVA payment. Unfortunately, it is generally not possible to tell which identifiers were created by the same bidder.

One may wonder whether we can address some of these problems by using a combinatorial auction mechanism other than the GVA. It has been shown that *any* mechanism that avoids these issues (revenue nonmonotonicity, collusion, and false-name bidding) must have some unnatural properties [Rastegari *et al.*, 2007]. Still, several false-name-proof combinatorial auction mechanisms have been designed [Yokoo *et al.*, 2001; Yokoo, 2003; Yokoo *et*

al., 2004, 2006]. It is also possible to make the use of multiple identifiers suboptimal by **verifying** the identities of some of the bidders [Conitzer, 2007].

11 Computationally efficient truthful combinatorial auctions

Another problem with the GVA is that it requires us to solve the winner determination problem to optimality. In fact, to compute the GVA payments, we need to solve up to n additional WDP instances: for each winning bidder, we need to solve the problem again with that bidder omitted. (It should be noted that at least in some settings, some of the computational work can be reused across the different instances [Hershberger and Suri, 2001].) An obvious idea is to not solve the WDP to optimality, but rather to use an approximation algorithm that returns a solution that is close to optimal. However, this effectively changes the mechanism, and there is no reason to think that desirable properties such as truthfulness and voluntary participation will continue to hold [Nisan and Ronen, 2001].

Let us consider the special case of single-minded bidders. One natural approximation algorithm is the following. Sort the bids (S_i, v_i) by $v_i/|S_i|$, the per-item value of the bid, in descending order. Then, consider the bids in this order, and accept any bid that can still be accepted. For example, suppose the bids are, in sorted order, $(\{a\}, 11), (\{b, c\}, 20), (\{a, d\}, 18), (\{a, c\}, 16), (\{c\}, 7), (\{d\}, 6)$. The algorithm will accept the first two bids; the next three bids can then no longer be accepted, because one of their items has already been allocated; finally, the last bid is accepted. The total value of this allocation is $11 + 20 + 6 = 37$ (which is less than the $20 + 18 = 38$ that could have been obtained by accepting just the second and third bid). Now, if we wish to calculate the first bidder's GVA payment using this approximation algorithm, we must remove this bid, and run the algorithm again. After removing the first bid, the algorithm will actually accept the bids $(\{b, c\}, 20)$ and $(\{a, d\}, 18)$, the optimal solution. Hence, the first bidder's (approximated)

GVA payment is $38 - (37 - 11) = 12$. This is more than the bidder's valuation! It follows that this approximation of the GVA mechanism does not satisfy voluntary participation (and hence it is also not truthful, because the bidder would have been better off bidding 0).

However, we *can* use this approximation algorithm for the WDP to obtain a truthful mechanism that satisfies voluntary participation: we just need to compute the payments somewhat differently [Lehmann *et al.*, 2002]. For each winning bid, consider the first bid in the sorted list that was forced out by this bid. The ratio $v_i/|S_i|$ for that bid is what the winning bid must pay per item. For example, in the above instance, the first bid forced out by the winning bid $(\{a\}, 11)$ is $(\{a, d\}, 18)$. Hence, the bid $(\{a\}, 11)$ pays $18/|\{a, b\}| = 9$ per item—and since it wins only one item (a), that means it pays 9. As for the winning bid $(\{b, c\}, 20)$, the first bid forced out by it is $(\{c\}, 7)$ (it is *not* $(\{a, c\}, 16)$, because this bid was already forced out by $(\{a\}, 11)$), so the bid pays $7/|\{c\}| = 7$ per item, and since it wins two items, it pays 14. The final winning bid, $(\{d\}, 6)$, forces no other bids out and hence pays 0. Given that bidders are single-minded, this mechanism satisfies voluntary participation and truthfulness (each winning bidder pays the lowest amount that they could have bid while still winning).

The above approximation algorithm for the WDP has a worst-case approximation ratio of m, the number of items. If we sort the bids by $v_i/\sqrt{|S_i|}$ instead, then the approximation ratio is improved to $\sqrt{m}$.

There is a significant body of work on computationally efficient truthful combinatorial auctions: see, for example, [Nisan and Ronen, 2000; Mu'alem and Nisan, 2002; Bartal *et al.*, 2003; Archer *et al.*, 2003; Dobzinski *et al.*, 2006; Bikhchandani *et al.*, 2006; Dobzinski and Nisan, 2007a,b].

12 Iterative combinatorial auctions and preference elicitation

In a perfect world, every bidder would have a valuation function that can be concisely expressed in the bidding language of the combinatorial auction. Unfortunately, in reality, this is often not the case. This does not mean that bidders usually submit extremely long bids (such as an XOR of $2^m - 1$ bundles): this is too impractical, not only because it requires the communication of an exponential amount of information, but also because *determining* one's valuation for a given bundle is generally a nontrivial task. Instead, bidders bid on a few bundles on which they think their bids will be competitive. But they may not know exactly on which bundles they would be competitive, and if they do not bid on the right bundles, this results in decreased economic welfare.

A potential remedy for this is to, during the auction, give the bidders feedback on how they are doing in the auction. This means that we must abandon the sealed-bid format, and instead consider **iterative** auction mechanisms. We have already seen some examples of iterative auctions in the single-item context, namely the English, Japanese, and Dutch auctions. For example, in an English auction, a bidder knows if she is currently winning, and if she is not, she can choose to raise her bid. The English and Japanese auctions are **ascending** auctions.

It turns out that we can also create ascending combinatorial auctions. As an example, let us study the **iBundle** ascending combinatorial auction [Parkes and Ungar, 2000]. This auction maintains, for each bidder i, for each bundle S, a current price $p_i(S)$. In each round of the auction, the bidder is supposed to choose the bundle(s) most attractive to her at her current prices, that is, the set $\arg\max_S v_i(S) - p_i(S)$, and bid $p_i(S)$ on these bundles. The exception is if $v_i(S) - p_i(S) < 0$ for every bundle S, in which case the bidder is supposed to drop out. If a bidder follows this strategy, she is said to **bid straightforwardly**. At the end of the round, the winner determination problem is solved with the submitted bids.

Then, for every active bidder i that is not winning anything, for every bundle S that she bid on, the price $p_i(S)$ is increased by some predetermined amount ϵ. Eventually, there will be a round where every active bidder wins something, and at this point the auction terminates with the current allocation and payments. This auction is known to have some nice properties: for example, if the bidders' valuations satisfy a condition known as **buyer submodularity**, then straightforward bidding is an **ex-post equilibrium**, and the GVA outcome results [Ausubel and Milgrom, 2002]. Informally, buyer submodularity means that the more bidders are already present, the less having additional bidders adds to the final allocation value. Strategies are said to be in ex-post equilibrium if for each bidder i, it is optimal to follow the strategy, assuming that the other bidders follow their strategies (but regardless of what the other bidders' valuations are). Numerous other iterative combinatorial auctions have been proposed (for an overview, see [Parkes, 2006]), and inherent limitations of this approach have also been studied [Blumrosen and Nisan, 2005].

More generally, the auctioneer can sequentially ask the bidders various **queries** about their valuation functions, until the auctioneer has enough information to determine the final outcome. If the final outcome is (always) the GVA outcome, then responding truthfully to the auctioneer is an ex-post equilibrium. This flexible query-based approach is generally referred to as **preference elicitation** [Conen and Sandholm, 2001]. Common queries include **value queries** ("What is your valuation for bundle S?") and **demand queries** ("Given these prices, which bundle(s) do you prefer?"). Demand queries can use either **item prices**, where the price for a bundle is the sum of the prices of the individual items, or **bundle prices**, where each bundle has a separate price (as we saw in the iBundle auction above). For some restricted classes of valuation functions, it has been shown that a polynomial number of queries suffices to learn a bidder's valuation function completely (if we assume that the valuation function lies in that class). For example, a valuation function can be elicited using a number of (bundle-price) demand queries that is polynomial in the length of the function's XOR-representation [Lahaie and Parkes, 2004], though in general

an exponential number of queries is required if only item prices (and value queries) are allowed [Blum *et al.*, 2004]. Various other results have been obtained on classes of valuation functions that can(not) be elicited using a polynomial number of queries [Zinkevich *et al.*, 2003; Santi *et al.*, 2004; Conitzer *et al.*, 2005; Blumrosen and Nisan, 2005; Lahaie *et al.*, 2005].

Without any restrictions on the valuation functions, negative results are known: for example, solving the winner determination problem in general requires an exponential amount of communication, regardless of what types of query are used [Nisan and Segal, 2005].

13   Additional topics

In this final section, we mention a few additional topics and provide some references for the interested reader.

There are several important variants of (combinatorial) auctions, including **(combinatorial) reverse auctions** and **(combinatorial) exchanges**. In a reverse auction, the auctioneer seeks to *buy* one or more items, and the bidders submit bids indicating how much they need to be compensated to provide the items. In an exchange, bidders can act as buyers as well as sellers. While these variants display some significant similarities to regular (forward) auctions, there are also important differences [Sandholm *et al.*, 2002].

Quite a few researchers have tried to generalize Myerson's expected-revenue maximizing auction to combinatorial auctions; this turns out to be surprisingly difficult [Avery and Hendershott, 2000; Armstrong, 2000; Conitzer and Sandholm, 2004; Likhodedov and Sandholm, 2004, 2005]. A different take on this problem is to design **competitive** auctions, which obtain a revenue that is within a factor of the revenue that can be obtained with a single sale price [Goldberg *et al.*, 2006].

A final important direction is the design of **online auctions**. "Online" here does not refer to Internet auctions; rather, it refers to settings in which the bidders arrive at and

depart from the auction over time, and allocation decisions must be made before all the bidders have arrived. More detail can be found in, for example, [Lavi and Nisan, 2000; Awerbuch *et al.*, 2003; Friedman and Parkes, 2003; Kleinberg and Leighton, 2003; Blum *et al.*, 2003; Hajiaghayi *et al.*, 2004, 2005; Bredin and Parkes, 2005; Blum *et al.*, 2006; Babaioff *et al.*, 2007; Parkes and Duong, 2007; Hajiaghayi *et al.*, 2007].

References

Aaron Archer, Christos H. Papadimitriou, Kunal Talwar, and Eva Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2003.

Mark Armstrong. Optimal multi-object auctions. *Review of Economic Studies*, 67:455–481, 2000.

Lawrence Ausubel and Paul Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1, 2002. No. 1, Article 1.

Lawrence M. Ausubel and Paul Milgrom. The lovely but lonely Vickrey auction. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 1. MIT Press, 2006.

Christopher Avery and Terrence Hendershott. Bundling and optimal auctions of multiple products. *Review of Economic Studies*, 67:483–497, 2000.

Baruch Awerbuch, Yossi Azar, and Adam Meyerson. Reducing truth-telling online mechanisms to online optimization. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*, 2003.

Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2007.

Martin J. Bailey. The demand revealing process: to distribute the surplus. *Public Choice*, 91:107–126, 1997.

Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi-unit combinatorial auctions. In *Theoretical Aspects of Rationality and Knowledge (TARK)*, Bloomington, Indiana, USA, 2003.

Sushil Bikhchandani, Shurojit Chatterji, Ron Lavi, Ahuva Mu'alem, Noam Nisan, and Arunava Sen. Weak monotonicity characterizes deterministic dominant strategy implementation. *Econometrica*, 74(4):1109–1132, 2006.

Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 202–204, 2003.

Avrim Blum, Jeffrey Jackson, Tuomas Sandholm, and Martin Zinkevich. Preference elicitation and query learning. *Journal of Machine Learning Research*, 5:649–667, 2004.

Avrim Blum, Tuomas Sandholm, and Martin Zinkevich. Online algorithms for market clearing. *Journal of the ACM*, 53(5):845–879, 2006.

Liad Blumrosen and Noam Nisan. On the computational power of iterative auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 29–43, Vancouver, Canada, 2005.

Jonathan Bredin and David Parkes. Models for truthful online double auctions. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 50–59, Edinburgh, UK, 2005.

Ruggiero Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 882–889, Hakodate, Japan, 2006.

Ed H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

Wolfram Conen and Tuomas Sandholm. Preference elicitation in combinatorial auctions: Extended abstract. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 256–259, Tampa, FL, USA, October 2001.

Vincent Conitzer and Tuomas Sandholm. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 132–141, New York, NY, USA, 2004.

Vincent Conitzer and Tuomas Sandholm. Failures of the VCG mechanism in combinatorial auctions and exchanges. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 521–528, Hakodate, Japan, 2006.

Vincent Conitzer, Jonathan Derryberry, and Tuomas Sandholm. Combinatorial auctions with structured item graphs. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 212–218, San Jose, CA, 2004.

Vincent Conitzer, Tuomas Sandholm, and Paolo Santi. Combinatorial auctions with k-wise dependent valuations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 248–254, Pittsburgh, PA, 2005.

Vincent Conitzer. Limited verification of identities to induce false-name-proofness. In *Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 102–111, Brussels, Belgium, 2007.

Shahar Dobzinski and Noam Nisan. Limitations of vcg-based mechanisms. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*, pages 338–344, 2007.

Shahar Dobzinski and Noam Nisan. Mechanisms for multi-unit auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 346–351, San Diego, CA, USA, 2007.

Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*, pages 644–652, 2006.

Eric Friedman and David Parkes. Pricing WiFi at Starbucks – Issues in online mechanism design. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 240–241, San Diego, CA, USA, 2003.

Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 548–553, Stockholm, Sweden, August 1999.

Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:587–602, 1973.

Andrew Goldberg, Jason Hartline, Anna Karlin, Mike Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 2006.

Georg Gottlob and Gianluigi Greco. On the complexity of combinatorial auctions: Structured item graphs and hypertree decompositions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 152–161, San Diego, CA, USA, 2007.

Daniel A. Graham and Robert C. Marshall. Collusive bidder behavior at single-object second-price and English auctions. *Journal of Political Economy*, 95(6):1217–1239, 1987.

Jerry Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.

Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

Mingyu Guo and Vincent Conitzer. Worst-case optimal redistribution of VCG payments. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 30–39, San Diego, CA, USA, 2007.

Mohammad T. Hajiaghayi, Robert Kleinberg, and David C. Parkes. Adaptive limited-supply online auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 71–80, New York, NY, USA, 2004.

Mohammad T. Hajiaghayi, Robert Kleinberg, Mohammad Mahdian, and David C. Parkes. Online auctions with re-usable goods. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 165–174, Vancouver, Canada, 2005.

Mohammad T. Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Vancouver, BC, Canada, 2007.

John Hershberger and Subhash Suri. Vickrey prices and shortest paths: What is an edge worth? In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, 2001.

Leonid Khachiyan. A polynomial algorithm in linear programming. *Soviet Math. Doklady*, 20:191–194, 1979.

Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, pages 594–605, 2003.

Sébastien Lahaie and David Parkes. Applying learning algorithms to preference elicitation. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 180–188, New York, NY, USA, 2004.

Sébastien Lahaie, Florin Constantin, and David C. Parkes. More on the power of demand queries in combinatorial auctions: Learning atomic languages and handling incentives. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 959–964, Edinburgh, UK, 2005.

Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 233–241, Minneapolis, MN, USA, 2000.

Daniel Lehmann, Lidian Ita O'Callaghan, and Yoav Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

Kevin Leyton-Brown, Yoav Shoham, and Moshe Tennenholtz. Bidding clubs: Institutionalized collusion in auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, Minneapolis, MN, USA, 2000.

Kevin Leyton-Brown, Yoav Shoham, and Moshe Tennenholtz. Bidding clubs in first-price auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 373–378, Edmonton, Canada, 2002.

Anton Likhodedov and Tuomas Sandholm. Methods for boosting revenue in combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 232–237, San Jose, CA, 2004.

Anton Likhodedov and Tuomas Sandholm. Approximating revenue-maximizing combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Pittsburgh, PA, 2005.

Hervé Moulin. Efficient, strategy-proof and almost budget-balanced assignment, March 2007. Working Paper.

Ahuva Mu'alem and Noam Nisan. Truthful approximate mechanisms for restricted combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 379–384, Edmonton, Canada, July 2002.

Roger Myerson. Incentive compatibility and the bargaining problem. *Econometrica*, 41(1), 1979.

Roger Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 242–252, Minneapolis, MN, USA, 2000.

Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 2005.

Noam Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 1–12, Minneapolis, MN, USA, 2000.

David C. Parkes and Quang Duong. An ironing-based approach to adaptive online mechanism design in single-valued domains. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Vancouver, BC, Canada, 2007.

David Parkes and Lyle Ungar. Iterative combinatorial auctions: Theory and practice. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 74–81, Austin, TX, August 2000.

David Parkes. Iterative combinatorial auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 3. MIT Press, 2006.

Michal Penn and Moshe Tennenholtz. Constrained multi-object auctions and *b*-matching. *Information Processing Letters*, 75(1–2):29–34, July 2000.

Ryan Porter, Yoav Shoham, and Moshe Tennenholtz. Fair imposition. *Journal of Economic Theory*, 118:209–228, 2004.

Baharak Rastegari, Anne Condon, and Kevin Leyton-Brown. Revenue monotonicity in combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Vancouver, BC, Canada, 2007.

John G. Riley and William F. Samuelson. Optimal auctions. *American Economic Review*, 71:381–392, June 1981.

Michael Rothkopf, Aleksandar Pekeč, and Ronald Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.

Tuomas Sandholm and Subhash Suri. BOB: Improved winner determination in combinatorial auctions and generalizations. *Artificial Intelligence*, 145:33–58, 2003.

Tuomas Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. Winner determination in combinatorial auction generalizations. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 69–76, Bologna, Italy, July 2002.

Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, January 2002.

Tuomas Sandholm. Optimal winner determination algorithms. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, pages 337–368. MIT Press, 2006.

Paolo Santi, Vincent Conitzer, and Tuomas Sandholm. Towards a characterization of polynomial preference elicitation with value queries in combinatorial auctions. In *Conference on Learning Theory (COLT)*, pages 1–16, Banff, Alberta, Canada, 2004.

Moshe Tennenholtz. Some tractable combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Austin, TX, August 2000.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

Makoto Yokoo, Yuko Sakurai, and Shigeo Matsubara. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence*, 130(2):167–181, 2001.

Makoto Yokoo, Yuko Sakurai, and Shigeo Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.

Makoto Yokoo, Toshihiro Matsutani, and Atsushi Iwasaki. False-name-proof combinatorial auction protocol: Groves mechanism with submodular approximation. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1135–1142, Hakodate, Japan, 2006.

Makoto Yokoo. The characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 733–742, Acapulco, Mexico, 2003.

Martin Zinkevich, Avrim Blum, and Tuomas Sandholm. On polynomial-time preference elicitation with value queries. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 176–185, San Diego, CA, USA, 2003.