Chapter 5

Difficulties for Classical Mechanism Design

The classical study of mechanism design has not directly concerned itself with the more computational questions of the process, such as the representations of the outcome and preference spaces, or the complexity of choosing the optimal outcome. Some of the computational implications of the classical mechanisms are clear and direct. For example, to execute the VCG mechanism, we typically need to solve the optimization problem instance with all agents included, as well as, for each agent, the instance where that agent is removed. (However, in some cases, the structure of the domain can be used to solve all these instances simultaneously with an asymptotic time complexity that is the same as that of a single optimization [Hershberger and Suri, 2001].) But the issues run deeper than that. When the setting is complex and computation is limited, the optimizations must be approximated. This effectively results in a different mechanism, which may no longer be truthful—and its strategic equilibria (*e.g.*, Nash equilibria) may be terrible even when the approximation algorithm *per se* is very good. The resulting challenge is to design special approximation algorithms that do motivate the agents to report their preferences truthfully. Viewed differently, the challenge is to design special truthful mechanisms whose outcomes are at least reasonably good, and can be computed efficiently. This line of research, which has been called *algorithmic mechanism design* [Nisan and Ronen, 2001], has produced a number of interesting results [Nisan and Ronen, 2001, 2000; Feigenbaum *et al.*, 2001; Lehmann *et al.*, 2002; Mu'alem and Nisan, 2002; Archer *et al.*, 2003; Bartal *et al.*, 2003].

There have been various other directions in the study of mechanism design from a computer science perspective. One direction close to algorithmic mechanism design is the design of *anytime* mechanisms, which produce better outcomes as they are given more time to compute, but nevertheless maintain good incentive properties [Parkes and Schoenebeck, 2004]. Another goal that has been pursued is *distributing* the mechanism's computation across the agents [Parkes and Shneidman, 2004; Brandt and Sandholm, 2004b,a, 2005b,c,a; Izmalkov *et al.*, 2005; Petcu *et al.*, 2006]. A different direction is the design of mechanisms in task-allocation settings where the agents may fail to accomplish the task, and the failure probabilities need to be elicited from the agents, as well as the costs [Porter *et al.*, 2002; Dash *et al.*, 2004]. Rather than specifying a complete mechanism, another approach that has been considered is to provide the agents with limited data from a centralized optimization, and let the agents work out the remainder of the transactions. Specifically, the optimal allocation is given, as well as some bounds on reasonable prices, in such a way that the agents do not have an incentive to misreport [Bartal *et al.*, 2004]. (This has the advantage of circumventing, or at least delaying until later, results such as the Myerson-Satterthwaite impossibility theorem mentioned previously.)

This chapter provides some results on mechanism design in expressive preference aggregation settings. Unlike some the results mentioned above, these results are not inherently computational: rather, they are pure mechanism design results that are driven by the expressive nature of the preference aggregation problems under study. (Of course, computational advances are what has made running mechanisms in such expressive domains possible.) In Section 5.1, we study two related vulnerabilities of the VCG (Clarke) mechanism in combinatorial auctions and exchanges: low revenue/high cost, and collusion. Specifically, it will show how much worse these vulnerabilities are in these settings than in single-item settings [Conitzer and Sandholm, 2006d]. In Section 5.2, we study mechanism design for expressive preference aggregation for donations to (charitable) causes [Conitzer and Sandholm, 2004e].

5.1 VCG failures in combinatorial auctions and exchanges

The VCG mechanism is the canonical payment scheme for motivating the bidders to bid truthfully in combinatorial auctions and exchanges; if the setting is general enough, under some requirements, it is the only one [Green and Laffont, 1977; Lavi *et al.*, 2003; Yokoo, 2003]. Unfortunately, there are also many problems with the VCG mechanism [Rothkopf *et al.*, 1990; Sandholm, 2000; Ausubel and Milgrom, 2006]. In this section, we discuss two related problems: the VCG mechanism is vulnerable to collusion, and may lead to low revenue/high payment for the auctioneer. It is well-known that these problems occur even in single-item auctions (where the VCG mechanism specializes to the Vickrey or second-price sealed-bid auction). However, in the single-item setting, these problems are not as severe. For example, in a Vickrey auction, it is not possible for colluders to obtain the item at a price less than the bid of any other bidder. Additionally, in a Vickrey auction, various types of revenue equivalence with (for example) first-price sealed-bid auctions hold. As we will show, in the multi-item setting these properties do not hold and can be violated to an arbitrary extent. Some isolated examples of such problems with the VCG mechanism in multi-item settings have already been noted in the literature [Ausubel and Milgrom, 2006; Yokoo *et al.*, 2004; Archer and Tardos, 2002] (these will be discussed later in the section). In contrast, our goal in this section is to give a *comprehensive* characterization of how severe these problems can be and when these severe problems can occur. For the various variants of combinatorial auctions and exchanges, we study the following single problem that relates both issues under consideration: *Given some of the bids, how bad can the remaining bidders make the outcome?* Informally, "bad" here means that the remaining bidders are paid an inordinately large amount, or pay an inordinately small amount, relative to the goods they receive and/or provide. This is closely related to the problem of making revenue guarantees to the auctioneer. But it is also the collusion problem, if we conceive of the remaining bidders as colluders. (The collusion problem can become more difficult if the collusion is required to be *self-enforcing*. A collusion is self-enforcing when none of the colluders have an incentive to unilaterally deviate from the collusion. We will also study how this extra requirement affects our results.)

As it turns out, the fundamental problem of deciding how bad the remaining bidders can make the outcome is often computationally hard. Computational hardness here is a double-edged sword. On the one hand, if the problem is hard, collusion may not occur (or to a lesser extent) because the colluders cannot find a beneficial collusion. On the other hand, if the problem is hard, it is difficult to make strong revenue guarantees to the auctioneer. Of course, in either case, the computational hardness may be overcome in practice if the stakes are high enough.

All the results in this section hold even when all bidders are *single-minded*, that is, they bid only on a single bundle of items. Hence, we do not need to discuss bidding languages.

5.1.1 Combinatorial (forward) auctions

We recall that in a *combinatorial auction*, there is a set of items $I = \{s_1, s_2, \ldots, s_m\}$ for sale. A bid takes the form $b = (B, v)$, where $B \subseteq I$ and $v \in \mathbb{R}$. The winner determination problem is to label bids as accepted or rejected, to maximize the sum of the values of the accepted bids, under the constraint that no item occurs in more than one accepted bid. (This is assuming *free disposal*: items do not have to be allocated to anyone.)

Motivating example

(A similar example to the one described in this subsubsection has been given before [Ausubel and Milgrom, 2006], and examples of vulnerability to false-name bidding in combinatorial auctions [Yokoo *et al.*, 2004] can in fact also be used to demonstrate the basic point. We include this subsubsection for completeness.) Consider an auction with two items, s_1 and s_2. Suppose we have collected two bids (from different bidders), both $(\{s_1, s_2\}, N)$. If these are the only two bids, one of the bidders will be awarded both the items and, under the VCG mechanism, will have to pay N. However, suppose two more bids (by different bidders) come in: $(\{s_1\}, N+1)$ and $(\{s_2\}, N+1)$. Then these bids will win. Moreover, neither winning bidder will have to pay anything! (This is because a winning bidder's item would simply be thrown away if that winning bidder were removed.)

This example demonstrates a number of issues. First, the addition of more bidders can actually decrease the auctioneer's revenue from an arbitrary amount to 0. Second, the VCG mechanism is not revenue-equivalent to the sealed-bid first-price mechanism in combinatorial auctions, even when all bidders' true valuations are common knowledge[1]—unlike in the single-item case. Third, even when the other bidders by themselves would generate nonnegative revenue for the auctioneer under the VCG mechanism, it is possible that two colluders can bid so as to receive all the items without paying anything.

[1]Consider the above example with $N \geq 9$ and suppose that the four bids reflect the bidders' true valuations—since bidding truthfully is a weakly dominant strategy in the VCG mechanism. Running a first-price sealed bid auction in this setting, when all bidders' valuations are common knowledge, will not generate expected revenue less than $\frac{N}{8}$. For suppose the expected revenue is less than this. Then the probability that the revenue is at least $\frac{N}{4}$ must be less than $\frac{1}{2}$ by Markov's inequality. So, bidding $(\{A, B\}, \frac{N}{4})$ will win any bidder both items with probability at least $\frac{1}{2}$, leading to an expected utility of at least $\frac{1}{2}(N - \frac{N}{4}) = \frac{3N}{8}$. Because at most one of the three bidders with valuations $(\{A\}, N+1)$ or $(\{A, B\}, N)$ can win its desired bundle, it follows that at least one of these bidders has a probability of at most $\frac{1}{3}$ of winning its desired bundle, and thus has an expected utility of at most $\frac{N+1}{3}$. Because $N \geq 9$, $\frac{3N}{8} > \frac{N+1}{3}$, so this bidder would be better off bidding $(\{A, B\}, \frac{N}{4})$—contradicting the assumption that we are in equilibrium.

The following sums up the properties of this example.

Proposition 5 *In a forward auction (even with only 2 items), the following can hold simultaneously: 1. The winning bidders pay nothing under the VCG mechanism; 2. If the winning bids are removed, the remaining bids generate revenue N under the VCG mechanism; 3. If these bids were truthful (as we would expect under VCG), then if we had run a first-price sealed-bid auction instead (and the bidders' valuations were common knowledge), any equilibrium would have generated revenue $\Theta(N)$.*

Characterization

We now characterize the settings where, given the noncolluders' bids, the colluders can receive all the items for free.

Lemma 9 *If the colluders receive all the items at cost 0, then for any positive bid on a bundle B of items by a noncolluder, at least two of the colluders receive an item from B.*

Proof: Suppose that for some positive bid b on a bundle B by a noncolluder i, one of the colluders c receives all the items in B (and possibly others). Then, in the auction where we remove that colluder's bids, one possible allocation gives every remaining bidder all the goods that bidder received in the original auction; additionally, it gives i all the items in bundle B; and it disposes of all the other items c received in the original auction. With this allocation, the total value of the accepted bids by bidders other than c is at least $v(b)$ more than in the original auction. Because the total value obtained in the new auction is at least the value of this particular allocation, it follows that c imposes a negative externality of at least $v(b)$ on the other bidders, and will pay at least $v(b)$. But this contradicts the fact that no colluder pays anything; and hence it follows that for any positive bid b on a bundle B by a noncolluder i, at least two of the colluders receive an item from B. ■

Lemma 10 *Suppose all the items in the auction can be divided among the colluders in such a way that, for any positive bid on a bundle of items B by a noncolluder, at least two of the colluders receive an item from B. Then the colluders can receive all the items at cost 0.*

Proof: For the given partition of items among the noncolluders, let each colluder place a bid with an extremely large value on the bundle consisting of the items assigned to him in the partition. (For instance, twice the sum of the values of all noncolluders' bids.) Then, the auction will clear awarding each colluder the items assigned to him by the partition. Moreover, if we remove the bids of one of the colluders, all the remaining colluders' bids will still win—and thus none of the noncollu ders' bids will win, because each such bid requires items assigned to at least two colluders by the partition (and at least one of them is still in the auction and wins th ese items). Thus, each colluder (individually) imposes no externality on the other bidders. ■

Combining these two lemmas, we get:

Theorem 24 *The colluders can receive all the items at cost 0 if and only if it is possible to divide the items among the colluders in such a way that, for any positive bid B by a noncolluder, at least two colluders receive an item from B.*

Self-enforcing collusion

It turns out that requiring that the collusion is self-enforcing (i.e., no colluder has an incentive to unilaterally deviate) is no harder for the colluders:

Theorem 25 *Whenever the colluders can receive all the items for free, they can also receive them all for free in a self-enforcing way.*

Proof: Let each colluder bid on the same bundle as before; but, increase the bid value of each colluder by an amount that exceeds the utility that any colluder can get from any bundle of items. The colluders will continue to receive all the items at a cost of 0. Now, the only reason that a colluder may wish to deviate from this is that the colluder wishes to obtain items outside of the colluder's assigned bundle. However, doing so would prevent one of the other bundles from being awarded to its designated colluder. This would cause a decrease in the total value of bids awarded to bidders other than the deviating colluder that exceeds the utility of the deviating colluder for any bundle, and the deviating colluder would have to pay for this decrease under the VCG mechanism. Therefore, there is no incentive for the colluder to deviate. ■

Complexity

In order to collude in the manner described above, the n colluders must solve the following computational problem.

Definition 21 (DIVIDE-SUBSETS) *Suppose we are given a set I, as well as a collection $R = \{S_1, \dots, S_q\}$ of subsets of it. We are asked whether I can be partitioned into n parts $T_1, T_2, \dots, T_n$ so that no subset $S_i \in R$ is contained in one of these parts.*

Theorem 26 *DIVIDE-SUBSETS is NP-complete, even when $n = 2$.*

Proof: The problem is technically identical to HYPERGRAPH-2-COLORABILITY, which is NP-complete [Garey and Johnson, 1979]. ■

This hardness result only states that it is hard to identify the *most* beneficial collusion, and one may wonder whether it is perhaps easier to find *some* beneficial collusion. It turns out that the hardness of the former problem implies the hardness of the latter problem: the utility functions of the colluders can always be such that only the most beneficial collusion actually benefits them, in which case the two problems are the same. This observation can also be applied to hardness results presented later in this section.

5.1.2 Combinatorial reverse auctions

We recall that in a *combinatorial reverse auction*, there is a set of items $I = \{s_1, s_2, \dots, s_m\}$ to be procured. A bid takes the form $b = (B, v)$, where $B \subseteq I$ and $v \in \mathbb{R}$. (Here, v represents the value that the bidder must be compensated by in order to provide the goods B.) The winner determination problem is to label bids as accepted or rejected, to minimize the sum of the values of the accepted bids, under the constraint that each item occurs in at least one accepted bid. (This is assuming free disposal.)

Motivating example

Consider a reverse auction with m items, $s_1, s_2, \ldots, s_m$. Suppose we have collected two bids (from different bidders), both $(\{s_1, s_2, \ldots, s_m\}, N)$. If these are the only two bids, one of the bidders will be chosen to provide all the goods, and, under the VCG mechanism, will be paid N. However, suppose m more bids (by different bidders) come in: $(\{s_1\}, 0), (\{s_2\}, 0), \ldots, (\{s_m\}, 0)$. Then, these m bids will win. Moreover, each bidder will be paid N under the VCG mechanism. (This is because without this bidder, we would have had to accept one of the original bids.) Thus, the total payment that needs to be made is mN.[2]

Again, this example demonstrates a number of issues. First, the addition of more bidders may actually increase the total amount that the auctioneer needs to pay. Second, the VCG mechanism requires much larger payments than a first-price auction in the case where all bidders' valuations are common knowledge. (The first-price mechanism will not require a total payment of more than N for these valuations in any pure-strategy equilibrium.[3]) Third, even when the other bidders by themselves would allow the auctioneer to procure the items at a low cost under the VCG mechanism, it is possible for m colluders to get paid m times as much for all the items.

The following sums up the properties of this example.

Proposition 6 *In a reverse auction, the following can hold simultaneously: 1. The winning bidders are paid mN under the VCG mechanism; 2. If the winning bids are removed, the remaining bids allow the auctioneer to procure everything at a cost of only N under the VCG mechanism; 3. If these bids were truthful (as we would expect under VCG), then if we had run a first-price sealed-bid reverse auction instead (and the bidders' valuations were common knowledge), any equilibrium in pure strategies would have required total payment of at most N. (However, there are also mixed-strategy equilibria with arbitrarily large expected tot al payment.)*

---

[2] Similar examples have been discovered in the context of purchasing paths in a graph [Archer and Tardos, 2002]. However, in that setting, the buyer does not seek to procure all of the items, and hence the examples cannot be applied directly to combinatorial reverse auctions.

[3] Consider the above example and suppose that the $n + 2$ bids reflect the bidders' true valuations—since bidding truthfully is a weakly dominant strategy in the VCG mechanism. Supposing that a pure-strategy equilibrium is being played, let the total payment to be made in this equilibrium be π. (We observe that the final allocation can still be uncertain, e.g. if there is a random tie-breaking rule.) Suppose $\pi > N$. Then, the expected utility for either one of the bidders interested in providing the whole bundle can never exceed $\pi - N$ (because the bidder will be paid 0 whenever none of its bids are accepted, and providing any items at all will cost it N). Moreover, it is not possible for both of these bidders to simultaneously have an expected utility of $\pi - N$ (as this would mean that both are paid π with certainty, contrary to the fact that the total payment is π). It follows at least one has an expected utility of $\pi - N - \epsilon$ for some $\epsilon > 0$. But then this bidder would be better off bidding $\pi - \frac{\epsilon}{2}$ for the whole bundle, which would be accepted with certainty and give an expected utility of $\pi - N - \frac{\epsilon}{2}$. It follows that the total payment in a pure-strategy equilibrium cannot exceed N. Perhaps surprisingly, the first-price combinatorial reverse auction for this example (with commonly known true valuations corresponding to the given bids) actually has mixed-strategy equilibria with arbitrarily high expected payments.

Characterization

Letting N be the sum of the values of the accepted bids when all the colluders' bids are taken out,[4] it is clear that no colluder can be paid more than N. (With the colluder's bid, the sum of the values of others' accepted bids is still at least 0; without it, it can be at most N, because in the worst case the auctioneer can accept the bids that would be accepted if none of the colluders are present.) In this subsubsection, we will identify a necessary and sufficient condition for the colluders to be able to each receive N.

Lemma 11 *If a colluder receives N, then the items that it has to provide cannot be covered by a subset of the noncolluders' bids with cost less than N.*

Proof: If they could be covered by such a set, we could simply accept this set of bids (including those that were accepted already) rather than the colluder's bid, and increase the total cost by less than N. Thus, the colluder's VCG payment is less than N. ■

Thus, in order for each of the n colluders to be able to receive N, it is necessary that there exist n disjoint subsets of the items, each of which cannot be covered with a subset of the noncolluders' bids with total value less than N. The next lemma shows that this condition is also sufficient.

Lemma 12 *If there are n disjoint sets of items $R_1, \ldots, R_n$, each of which cannot be covered by a subset of the noncolluders' bids with cost less than N, then n colluders can be paid N each.*

Proof: Let colluder i (for $i < n$) bid $(R_i, 0)$, and let colluder n bid $(R_n \cup (S - \bigcup_i R_i), 0)$. Then the total cost of all accepted bids with all the colluders is 0; but when one colluder is omitted, the items it won cannot be covered at a cost less than N (because its bid contained one of the R_i). Thus, each colluder's VCG payment is N. ■

The next lemma shows that the necessary and sufficient condition above is equivalent to being able to *partition* all the items into n sets, so that no element of the partition can be covered by a subset of the noncolluders' bids with total value less than N. That is, we can restrict our attention to the case where the subsets exhaust all the items.

Lemma 13 *The condition of Lemma 12 is satisfied if and only if it is possible to partition the items into $T_1, \ldots, T_n$ such that no T_i can be covered by a subset of the noncolluders' bids with cost less than N.*

Proof: The "if" part is trivial: given T_i that satisfy the condition of this lemma, simply let $R_i = T_i$. For the "only if" part, given R_i that satisfy the condition of Lemma 12, let $T_i = R_i$ for $i < n$, and $T_n = R_n \cup (S - \bigcup_i R_i)$. We observe that this last set can also not be covered at a cost of less than N because it contains R_n. ■

Combining all the lemmas, we get:

[4] We assume, as is commonly done in settings such as these, that a feasible solution still exists when all the colluders' bids are removed.

Theorem 27 *The n colluders can receive a payment of N each (simultaneously), where N is the sum of the values of the accepted bids when all the colluders' bids are removed, if and only if it is possible to partition the items into $T_1, \ldots, T_n$ such that no T_i can be covered by a subset of the noncolluders' bids with cost less than N.*

Self-enforcing collusion

Unlike the case of combinatorial forward auctions, in reverse auctions, a stronger condition is required if the collusion is also required to be self-enforcing.

Theorem 28 *The n colluders can receive a payment of N each (simultaneously), where N is the sum of the values of the accepted bids when all the colluders' bids are removed, if and only if it is possible to partition the items into $T_1, \ldots, T_n$ such that 1) no T_i can be covered by a subset of the noncolluders' bids with cost less than N, 2) for no colluder i, the following holds: there exists a subset $T_i' \subseteq T_i$ such that T_i' can be covered by a set of noncolluders' bids with total cost less than $v_i(T_i) - v_i(T_i - T_i')$ (the marginal savings to colluder i of not having to provide T_i').*

Proof: For the "if" part, each colluder i can bid on T_i with a value of 0. As in the above, this will give each colluder a payment of N. Moreover, no colluder i has an incentive to deviate, for the following reasons. Under the VCG mechanism, it is not possible to change a bidder i's bid in such a way that the allocation to i remains the same, but the payment to i changes. Therefore, we only need to consider what happens if colluder i bids on a different bundle. Bidding on items outside T_i cannot increase the payment to i because the other colluders are bidding on these items with a value of 0. Therefore, the only deviation that can possibly be advantageous is to bid on a subset T_i'' of T_i. Let $T_i' = T_i - T_i''$. If the colluder bids on T_i'' (with, say, value 0), then the payment to colluder i will decrease by the total cost of covering T_i' with noncolluder bids. By the assumption in the theorem, this total cost is at least $v_i(T_i) - v_i(T_i'')$, the marginal savings to colluder i of not having to provide T_i'. It follows that the bid does not make the colluder better off.

For the "only if" part, we already know by Theorem 27 that in order for the n colluders to receive a payment of N each (simultaneously), it must be possible to partition the items into $T_1, \ldots, T_n$ such that no T_i can be covered by a subset of the noncolluders' bids with cost less than N (so that colluder i can bid on T_i with a value of 0 to achieve the desired outcome). But if for some colluder i, there exists a subset $T_i' \subseteq T_i$ such that T_i' can be covered by a set of noncolluders' bids with total cost less than $v_i(T_i) - v_i(T_i - T_i')$, then this colluder would be better off bidding a value of 0 for $T_i - T_i'$ instead, because this would decrease the payment to colluder i by less than the marginal savings to colluder i of not having to provide T_i'. Hence the collusion would not be self-enforcing. ∎

Complexity

In order to collude in the manner described above, the n colluders must solve the following computational problem.

Definition 22 (CRITICAL-PARTITION) *We are given a set of items I, a collection of bids (S_i, v_i) where $S_i \subseteq I$ and $v_i \in \mathbb{R}$, and a number n. Say that the* cost *of a subset of these bids is the sum of* *their v_i; and that the cost $c(T)$ of a subset $T \subseteq I$ is the lowest cost of any subset of the bids whose S_i cover T. We are asked whether there exists a partition of I into n disjoint subsets $T_1, T_2, \ldots, T_n$, such that for any $1 \leq i \leq n$, $c(T_i) = c(I)$.*

Theorem 29 *Even when the bids are so that a partition $T_1, \ldots, T_n$ is a solution if and only if no set $I - T_i$ covers all items in a bid, CRITICAL-PARTITION is NP-complete (even with $n = 2$).*

Proof: The problem is in NP in this case because given a partition $T_1, \ldots, T_n$, it is easy to check if any set $I - T_i$ covers all items in a bid.

To show NP-hardness, we reduce an arbitrary NAESAT[5] instance (given by a set of clauses C over a set of variables V, with each variable occurring at most once in any clause) to the following CRITICAL-PARTITION instance with $n = 2$ (where we are trying to partition into T_1 and T_2). Let I be as follows. For every variable $v \in V$, there are two items labeled s_{+v} and s_{-v}. Let the bids be as follows. For every variable $v \in V$, there is a bid $(\{s_{+v}, s_{-v}\}, 2)$. For every clause $c \in C$, there are two bids $(\{s_l : l \in c\}, 2m_c - 1)$ and $(\{s_l : -l \in c\}, 2m_c - 1)$ where m_c is the number of literals occurring in c.

First we show that this instance satisfies the condition that a partition $T_1, \ldots, T_n$ is a solution if and only if no set $I - T_i$ covers all items in a bid. First, we observe that $c(I) = |I|$ (we can use all the bids of the form $(\{s_{+v}, s_{-v}\}, 2)$, getting a per-item cost of 1; no other bid gives a lower per-item cost).

Now, if some set $I - T_i$ covers all the items in a bid of the form $(\{s_{+v}, s_{-v}\}, 2)$, then $c(T_i) \leq 2|I| - 2$ (because we can simply omit this bid from the solution for all the items). If some set $I - T_i$ covers all the items in a bid of the form $(\{s_l : l \in c\}, 2m_c - 1)$, then $c(T_i) = |I| - 1$. (This is because we can now accept the "complement" bid $(\{s_l : -l \in c\}, 2m_c - 1)$, and we will have covered all the items s_{+v} and s_{-v} in T_i such that v occurs in c (precisely $2m_c$ items, because variables do not reoccur within a clause); for any other item s_{+v} or s_{-v}, we can accept the bid $(\{s_{+v}, s_{-v}\}, 2)$, and we need to accept at most $|V| - m_c$ such bids, leading to a total cost of $2m_c - 1 + 2(|V| - m_c) = |I| - 1$.)

On the other hand, suppose there is no set $I - T_i$ that covers all the items in a bid. Then, either T_i must include at precisely one of s_v and s_{-v}. (Otherwise one T_i would include neither and $I - T_i$ would cover all items in the bid $(\{s_{+v}, s_{-v}\}, 2)$.) Thus, when we are trying to cover T_i, covering items in it with bids of the form $(\{s_{+v}, s_{-v}\}, 2)$ would result in a per-item cost of 2. On the other hand, covering items in it with bids of the form $(\{s_l : l \in c\}, 2m_c - 1)$ or $(\{s_l : -l \in c\}, 2m_c - 1)$ would result in a per-item cost of at least $\frac{2m_c - 1}{m_c - 1} > 2$ (because at most $m_c - 1$ of the m_c items in the bid can be in T_i, otherwise T_i would cover all the items in the bid; but $T_i = I - T_{3-i}$ which by assumption does not cover all the items in any bid). It follows that $C(T_i) = 2|V| = |I| = c(I)$.

Now we show that the two instances are equivalent. First suppose there exists a solution to the NAESAT instance. Then partition the elements as $T_1 = \{s_l : l = \textit{true}\}$ and $T_2 = \{s_l : l = \textit{false}\}$, according to this solution. Clearly neither of $I - T_i = T_{3-i}$ covers a bid of the form $(\{s_{+v}, s_{-v}\}, 2)$. Also, because no clause has all its literals set to the same value (we have a NAESAT solution), the items in a corresponding bid $(\{s_l : l \in c\}, 2m_c - 1)$ or $(\{s_l : -l \in c\}, 2m_c - 1)$ are not all in

[5]The goal in NAESAT is to assign truth values to all variables in such a way that there is no clause with all its literals set to *true*, and no clause with all its literals set to *false*.

the same set. By the previously proved property, it follows that this partition is a solution to the CRITICAL-PARTITION instance.

On the other hand, suppose that there exists a solution to the CRITICAL-PARTITION instance. Then label a literal *true* if $s_l \in T_1$, and *false* otherwise. By the previously proved property, because $(\{s_{+v}, s_{-v}\}, 2)$ is a bid, only one of s_{+v} and s_{-v} can be in $T_1 = I - T_2$, so this provides a consistent setting of the literals. Additionally, because $(\{s_l : l \in c\}, 2m_c - 1)$ is a bid, not all the s_l in that bid can be in $T_1 = I - T_2$. It follows that some of the literals $l \in c$ are set to *false*. Similarly, not all the s_l in that bid can be in $T_2 = I - T_1$, so some of the literals $l \in c$ are set to *true*. It follows that this assignment of truth values to variables is a solution to the NAESAT instance. ∎

5.1.3 Combinatorial forward (or reverse) auctions without free disposal

We recall that a combinatorial forward auction *without free disposal* is exactly the same as one with free disposal, with the exception that every item must be allocated to some bidder. Recall from Section 2.2 that since we are looking for an exact cover of the items, and negative bids may be of use, combinatorial forward auctions are technically identical to combinatorial reverse auctions.

Motivating example

Consider a forward auction with two nondisposable items, s_1 and s_2. Suppose we have collected two bids (from different bidders), both $(\{s_1, s_2\}, N)$. If these are the only two bids, one of the bidders will be awarded both the items and, under the VCG mechanism, will have to pay N. However, suppose two more bids (by different bidders) come in: $(\{s_1\}, N + M)$ and $(\{s_2\}, N + M)$, with $M > 0$. Then these bids will win. Moreover, because without free disposal, we cannot accept either of these bids without the other, each of these bidders will be *paid* M under the VCG mechanism!

Again, this example demonstrates a number of issues. First, additional bidders may change the auctioneer's revenue from an arbitrarily large positive amount to an arbitrarily large negative amount (an arbitrarily large cost). Second, the VCG mechanism may require arbitrarily large payments from the auctioneer even in cases where a first-price auction would actually generate revenue for the auctioneer, in the case where all bidders' valuations are common knowledge. (The first-price mechanism will generate a revenue of at least N for these valuations in any pure-strategy equilibrium.[6]) Third, even when the other bidders by themselves would generate positive revenue for the

[6] Consider the above example and suppose that the four bids reflect the bidders' true valuations—since bidding truthfully is a weakly dominant strategy in the VCG mechanism. Supposing that a pure-strategy equilibrium is being played, let the total revenue to the auctioneer be ρ, where ρ is possibly negative. (We observe that the final allocation can still be uncertain, e.g. if there is a random tie-breaking rule.) Suppose $\rho < N$. Then the expected utility for either of the bidders interested in providing the whole bundle is at most $N - \rho$. (If the bidder receives a singleton item, its utility is $-\infty$; if it receives nothing, its utility is 0; if it receives both items, its utility is $N - \rho$.) Moreover, it is not possible for both of these bidders to both have an expected utility of $N - \rho$, as this would mean they both receive both items with probability 1. It follows that at least one of them has an expected utility of $N - \rho - \epsilon$ where $\epsilon > 0$. But then this bidder would be better off bidding $\rho + \frac{\epsilon}{2}$, as this bid would be accepted with certainty and give an expected utility of $N - \rho - \frac{\epsilon}{2}$. It follows that the expected revenue in a pure-strategy equilibrium cannot be less than N. Similarly to the case of the combinatorial reverse auction with free disposal, there are mixed-strategy equilibria in the first-price auction where the auctioneer is forced to make arbitrarily large payments.

auctioneer under the VCG mechanism, it is possible that two colluders can make the auctioneer pay each of them an arbitrarily large amount.

The following sums up the properties of this example.

Proposition 7 *In a forward auction without free disposal (even with only two items), the following can hold simultaneously: 1. Each winning bidder is paid an arbitrary amount M under the VCG mechanism (where M depends only on the winners' bids); 2. If the winning bids are removed, the remaining bids actually generate revenue N to the auctioneer under the VCG mechanism; 3. If these bids were truthful (as we would expect under VCG), then if we had run a first-price sealed-bid auction instead (and the bidders' valuations were common knowledge), any equilibrium in pure strategies would have generated revenue N. (However, there are mixed-strategy equilibria with arbitrarily large cost to the auctioneer.)*

Characterization

In this subsubsection, we will identify a necessary and sufficient condition for the colluders to be able to each receive an arbitrary amount. Let $v(b)$ denote the value of bid b.

Lemma 14 *If each colluder receives a payment of more than $2\sum\limits_{d}|v(b_d)|$ (where d ranges over the noncolluders), then for each colluder c, the set of all items awarded to either that colluder or a noncolluder (that is, $s_c \cup \bigcup_d s_d$, where s_b is the set of items awarded to bidder b and d ranges over the noncolluders) cannot be covered exactly with bids from the noncolluders.*

Proof: Say that the sum of the values of accepted noncolluder bids is D (which may be negative). Suppose that for one colluder c, the set of all items awarded to either her or a noncolluder (that is, $s_c \cup \bigcup_d s_d$) can be covered by a set of noncolluder bids of combined value C (which may be negative). Then removing colluder c can make the allocation at most $D - C$ worse to the other bidders (relative to their reported valuations), because we could simply accept the bids of combined value C and no longer accept the bids of combined value D, and keep the rest of the allocation the same. Thus, under VCG, that colluder should be rewarded at most $D - C \leq 2\sum\limits_{d}|v(b_d)|$. ■

Thus, in order for each colluder to be able to receive an arbitrarily large payment, it is necessary that there are n disjoint subsets of the items such that no such subset taken together with the remaining items can be covered exactly by the noncolluders' bids. Also, the set of remaining items must be exactly coverable by the noncolluders' bids (otherwise we cannot accept all the colluders' bids). The next lemma shows that this condition is also sufficient.

Lemma 15 *If it is possible to partition the items into $R_1, \ldots, R_n, R_{n+1}$ such that for no $1 \leq i \leq n$, $R_i \cup R_{n+1}$ can be covered exactly with bids from the noncolluders; and such that R_{n+1} can be covered exactly with bids from the noncolluders; then for any $M > 0$, n colluders can place additional bids such that each of them receives at least M.*

Proof: Let colluder i place a bid $(R_i, M + 3\sum\limits_{d}|v(b_d)|)$ (where d ranges over the noncolluders). All these bids will be accepted, because it is possible to do so by also accepting the noncolluder bids that cover R_{n+1} exactly; and these noncolluder bids will have a combined value of at least $-\sum_d |v(b_d)|$, so that the sum of the values of all accepted bids is at least $(3n-1)\sum_d |v(b_d)| + nM$. (We observe that if we do not accept all of the colluder bids, the sum of the values of all accepted bids is at most $(3(n-1)+1)\sum_d |v(b_d)| + (n-1)M = (3n-2)\sum_d |v(b_d)| + (n-1)M$, which is less.) Now, if the bid of colluder i is removed, it is no longer possible to accept all the remaining $n-1$ colluder bids, because $R_i \cup R_{n+1}$ cannot be covered exactly with noncolluder bids. It follows that the total value of all accepted bids when i's bid is removed can be at most $(3(n-2)+1)\sum_d |v(b_d)| + (n-2)M$. When i's bid is not omitted, the sum of the values of all accepted bids other than i's is at least $(3(n-1)-1)\sum_d |v(b_d)| + (n-1)M$. Subtracting the former quantity from this, we get that the VCG payment to i is at least $\sum_d |v(b_d)| + M$. ■

The next lemma shows that the necessary and sufficient condition above is equivalent to being able to partition all the items into n sets, so that no element of the partition can be covered exactly by a subset of the noncolluders' bids. That is, we can restrict our attention to the case where $R_{n+1} = \emptyset$.

Lemma 16 *The condition of Lemma 15 is satisfied if and only if the items can be partitioned into $T_1, \ldots, T_n$ such that no T_i can be covered exactly with bids from the noncolluders.*

Proof: For the "if" part: given T_i that satisfy the condition of this lemma, let $R_i = T_i$ for $i \leq n$, and $R_{n+1} = \emptyset$. Then no $R_i \cup R_{n+1} = T_i$ can be covered exactly with bids from the noncolluders, and $R_{n+1} = \emptyset$ can trivially be covered exactly with noncolluder bids. For the "only if" part: given R_i that satisfy the condition of Lemma 15, let $T_i = R_i$ for $i < n$, and let $T_n = R_n \cup R_{n+1}$. That T_n cannot be covered exactly by noncolluder bids now follows directly from the conditions of Lemma 15. But also, no T_i with $i < n$ can be covered exactly: because if it could, then we could cover $R_i \cup R_{n+1} = T_i \cup R_{n+1}$ using the bids that cover T_i exactly together with the bids that cover R_{n+1} exactly (which exist by the conditions of Lemma 15). ■

Combining all the lemmas, we get:

Theorem 30 *The n colluders can receive a payment of at least M each (simultaneously), where M is an arbitrarily large number, if and only if it is possible to partition the items into $T_1, \ldots, T_n$ such that no T_i can be covered exactly with bids from the noncolluders.*

Self-enforcing collusion

Again, a stronger condition is required if the collusion is also required to be self-enforcing.

Theorem 31 *The n colluders can receive a payment of at least M each (simultaneously), where M is an arbitrarily large number, if and only if it is possible to partition the items into $T_1, \ldots, T_n$ such that 1) no T_i can be covered exactly with noncolluder bids, 2) for no colluder i, the following holds: there exists a subset $T_i' \subseteq T_i$ such that T_i' can be covered exactly by a set of noncolluders' bids with total value greater than $v_i(T_i) - v_i(T_i - T_i')$ (the marginal value to colluder i of receiving T_i').*

Proof: For the "if" part, each colluder i can bid on T_i with a sufficiently large value. As in the above, this will give each colluder a payment of at least M. Moreover, no colluder i has an incentive to deviate, for the following reasons. Under the VCG mechanism, it is not possible to change a bidder i's bid in such a way that the allocation to i remains the same, but the payment to i changes. Therefore, we only need to consider what happens if colluder i bids on a different bundle. Bidding on items outside T_i will prevent one of the other colluders' bids from being accepted, leading to a severe reduction in the total value of the allocation, and therefore to a severe reduction in the payment to colluder i. Therefore, the only deviation that can possibly be advantageous is to bid on a subset T_i'' of T_i. Let $T_i' = T_i - T_i''$. If the colluder bids on T_i'' (with a sufficiently large value), one of two things may happen. First, it can be the case that it is not possible to exactly cover T_i' with noncolluder bids. If so, then it must be the case that one of the other colluders' bids cannot be accepted, leading again to a severe reduction in the payment to colluder i. Second, it can be the case that it is possible to exactly cover T_i' with noncolluder bids. In this case, the payment to colluder i will increase by the total value of this cover of T_i'. By the assumption in the theorem, this total value is at most $v_i(T_i) - v_i(T_i'')$, the marginal value to colluder i of receiving T_i'. It follows that the bid does not make the colluder better off.

For the "only if" part, we already know by Theorem 30 that in order for the n colluders to receive an arbitrarily large payment of at least M each (simultaneously), it must be possible to partition the items into $T_1, \ldots, T_n$ such that no T_i can be covered exactly with noncolluder bids (so that colluder i can bid on T_i with a sufficiently large value to achieve the desired outcome). But if for some colluder i, there exists a subset $T_i' \subseteq T_i$ such that T_i' can be covered exactly by a set of noncolluders' bids with total value greater than $v_i(T_i) - v_i(T_i - T_i')$, then this colluder would be better off bidding a sufficiently large value for $T_i - T_i'$ instead, because this would increase the payment to colluder i by more than the marginal value to colluder i of receiving T_i' as well. Hence the collusion would not be self-enforcing. ■

Complexity

In order to collude in the manner described above, the n colluders must solve the following computational problem.

Definition 23 (COVERLESS-PARTITION) *We are given a set I and a collection of subsets $S_1, S_2, \ldots, S_q \subseteq I$. We are asked whether there is a partition of I into subsets $T_1, T_2, \ldots T_n \subseteq I$ such that no T_i can be covered exactly by some of the S_i.*

Theorem 32 *Even if there is a singleton S_i for all but two elements a and b, and $n = 2$, COVERLESS-PARTITION is NP-complete.*

Proof: The problem is in NP in this case because given a partition T_1, T_2, either one of the T_i contains both a and b, in which case the other can be covered exactly with singleton sets; or they each contain one of a and b (say T_a contains a and T_b contains b). In the latter case, there is a cover of T_s if and only if it contains a subset containing s (the other elements in T_s can be covered with singleton sets), which can be checked in polynomial time.

To show that the problem is NP-hard, we reduce from SAT. Given an arbitrary SAT instance (given by a set of clauses C over variables V), let S be as follows. It contains a and b; for each variable $v \in V$, it contains an element s_v; and for each clause $c \in C$, it contains an element s_c. Let the collection of subsets be as follows. For every s_v, there is a subset $S_v = \{s_v\}$. For every s_c, there is a subset $S_c = \{s_c\}$. Finally, for every clause $c \in C$, there are two more subsets: one consisting of b, s_c, and all the variables that occur positively in c ($S_{c+} = \{b, s_c\} \cup \{s_v : +v \in c\}$), and one consisting of a, s_c, and all the variables that occur negatively in c ($S_{c-} = \{a, s_c\} \cup \{s_v : -v \in c\}$). We now show that the instances are equivalent.

First suppose there is a solution to the SAT instance, given by a labeling $t : V \rightarrow \{$*true, false*$\}$. Then let $\{a\} \cup \{s_v : t(v) =$*true*$\} \subseteq T_1$ and $\{b\} \cup \{s_v : t(v) =$*false*$\} \subseteq T_2$. Furthermore, if one of the variables occurring positively in c is set to *true*, let $s_c \in T_2$; otherwise, let $s_c \in T_1$. First, we claim that no subset S_{c+} is contained in some T_i. It is not contained in T_1 because it has b in it. If c is satisfied because of one of the variables v occurring positively in c is set to *true*, then $s_v \in T_1$, and because $s_v \in S_{c+}$, S_{c+} is not contained in T_2. Otherwise, $s_c \in T_1$, and again S_{c+} is not contained in T_2. Next, we claim that no subset S_{c-} is contained in some T_i. It is not contained in T_2 because it has a in it. If c is satisfied because of one of the variables v occurring positively in c is set to *true*, $s_c \in T_2$, and S_{c-} is not contained in T_1. Otherwise, one of the variables v occurring negatively in c must be set to *false*, so $v \in T_2$, and because $s_v \in S_{c-}$, again S_{c-} is not contained in T_1. Because only bids of the form S_{c+} or S_{c-} contain a or b, it follows that there is no exact cover of either T_1 or T_2, and we have a solution to the COVERLESS-PARTITION instance.

Now suppose there is a solution to the COVERLESS-PARTITION instance, given by a partition T_1, T_2. Because a and b cannot occur in the same T_i, suppose without loss of generality that $a \in T_1$ and $b \in T_2$. Then, set v to *true* if $s_v \in T_1$, and to *false* otherwise. Suppose that a given clause c is not satisfied with this assignment. This means that for all variables v that occur positively in c, $s_v \in T_2$, and for all variables v that occur negatively in c, $s_v \in T_1$. If $s_c \in T_1$, then S_{c-} is contained in T_1; and thus we can cover T_1 exactly with this set and singleton sets for the remaining elements. On the other hand, if $s_c \in T_2$, then S_{c+} is contained in T_2; and thus we can cover T_2 exactly with this set and singleton sets for the remaining elements. It follows that all clauses are satisfied with this assignment, and we have a solution to the SAT instance. ∎

An easier collusion problem

So far in this subsection, we have formulated the collusion problem so that *each* colluder should receive M, where M is an arbitrary amount. An easier problem for the colluders is to make sure that *together*, they receive M, where M is an arbitrary amount. Such a collusion may be less stable (because some of the colluders may be receiving very little). Nevertheless, as we will show, this type of collusion is possible whenever a weak (and easily verified, given the noncolluders' bids) condition holds: at least one item has no singleton bid on it. (A singleton bid is a bid on only one item.) We first show that this condition is necessary.

Lemma 17 *If at least one colluder receives a payment of more than* $\sum_d |v(b_d)|$ *(where d ranges over the noncolluders), then there is at least one item s on which no noncolluder places a singleton bid.*

Proof: If each item has a singleton noncolluder bid placed on it, then when we remove a colluder's bid, we can simply cover all the items in it with singleton bids (with a combined value of at least $-\sum_d |v(b_d)|$), and leave the rest of the allocation unchanged. It follows that the VCG payment to the colluder can be at most $\sum_d |v(b_d)|$). ■

We now show that the condition is sufficient.

Lemma 18 *If there is at least one item s on which no noncolluder places a singleton bid, then if one colluder bids $(\{s\}, 0)$, and the other colluder bids $(I - \{s\}, M + 2\sum_d |v(b_d)|)$ (for $M > 0$), then the total payment to the colluders is at least M.*

Proof: The colluders' bids will be the only accepted ones (because colluder 2's bid has a greater value than all other bids combined). If we removed colluder 2's bid, the total value of the accepted bids would be at most $\sum_d |v(b_d)|$), so colluder 2 will pay at most this much under the VCG mechanism. If we removed colluder 1's bid, colluder 2's bid could no longer be accepted (because $\{s\}$ cannot be covered by itself), and thus the total value of the accepted bids could be at most $\sum_d |v(b_d)|$). It follows that colluder 1 is paid at least $M + \sum_d |v(b_d)|$). So the total payment to the colluders is at least M ■

Combining the two lemmas, we get the desired result:

Theorem 33 *Two (or more) colluders can receive a total payment of M, where M is an arbitrarily large number, if and only if there is at least one item that has no singleton bid placed on it by a noncolluder.*

5.1.4 Combinatorial exchanges

We recall that in a *combinatorial exchange*, there is a set of items $I = \{s_1, s_2, \ldots, s_m\}$ that can be traded. A bid takes the form $b = (\lambda_1, \ldots, \lambda_m, v)$, where $\lambda_1, \ldots, \lambda_m, v \in \mathbb{R}$ (possibly negative). (Each λ_i is the number of units of the ith item that the bidder seeks to procure, and v is how much the bidder is wi lling to pay.) The winner determination problem is to label bids as accepted or rejected, under the constraint that the sum of the accepted vectors has its first m entries ≤ 0, to maximize the last entry of the sum of the accepted vectors. (This is assuming free disposal.) We will also use the notation $(\{(s_{i_1}, \lambda_{i_1}), (s_{i_2}, \lambda_{i_2}), \ldots, (s_{i_k}, \lambda_{i_k})\}, v)$ for representing a bid in which λ_{i_j} units of item s_{i_j} are demanded (and 0 units of each item that is not mentioned).

Characterization

In a combinatorial exchange with at least two items s_1 and s_2, let q_1 (respectively, q_2) be the total number of units of s_1 (respectively, s_2) offered for sale in bids so far (by noncolluders). Now consider the following two bids (by colluders): $(\{(s_1, q_1 + 1), (s_2, -q_2 - 1)\}, M + \sum_d |v(b_d)|)$ and $(\{(s_1, -q_1 - 1), (s_2, q_2 + 1)\}, M + \sum_d |v(b_d)|)$, where $M > 0$ and d ranges over the original (noncolluding) bids. Both these bids will be accepted (for otherwise, the total value of the accepted bids could be at most $M + 2\sum_d |v(b_d)| < 2(M + \sum_d |v(b_d)|)$). Moreover, if we remove one of these two bids, the other cannot be accepted (because its demand cannot be met), so the total value of the accepted bids can be at most $\sum_d |v(b_d)|$). It follows that the VCG payment to each of these two bidders is at least M. This proves the following:

Theorem 34 *In a combinatorial exchange with at least two items, for any set of bids by noncolluders, two colluders can place bids so that each of them will receive at least M, where M is an arbitrary amount. Moreover, each one receives exactly the items that the other provides, so that their net contribution in terms of items is nothing.*

This concludes the part of this dissertation analyzing problematic outcomes of the VCG mechanism in combinatorial auctions and exchanges. We will return to combinatorial auctions briefly in the next chapter, Section 6.5. In the next section, we consider mechanism design for negotiating over donations to charities.

5.2 Mechanism design for donations to charities

In this section, we study mechanism design for the setting of expressive preference aggregation for donations to charities described in Section 2.3. The rules that we described in that section for deciding on outcomes turn out not to be strategy-proof, as we will see shortly. This is not too surprising, because the mechanism described so far is, in a sense, a *first-price* mechanism, where the mechanism will extract as much payment from a bidder as her bid allows; and such mechanisms are typically not strategy-proof. In this section, we consider changing the rules to make bidding truthfully strategically optimal.

5.2.1 Strategic bids under the first-price mechanism

We first point out some reasons for bidders to misreport their preferences under the first-price mechanism described up to this point. First of all, even when there is only one charity, it may make sense to underbid one's true valuation for the charity. For example, suppose a bidder would like a charity to receive a certain amount x, but does not care if the charity receives more than that. Additionally, suppose that the other bids guarantee that the charity will receive at least x no matter what bid the bidder submits (and the bidder knows this). Then the bidder is best off not bidding at all (or submitting a utility for the charity of 0), to avoid having to make any payment. (This is an instance of the *free rider* problem [Mas-Colell *et al.*, 1995].)

With multiple charities, another kind of manipulation may occur, where the bidder attempts to steer others' payments towards her preferred charity. Suppose that there are two charities, and three bidders. The first bidder bids $u_1^1(\pi_{c_1}) = 1$ if $\pi_{c_1} \geq 1$, $u_1^1(\pi_{c_1}) = 0$ otherwise; $u_1^2(\pi_{c_2}) = 1$ if $\pi_{c_2} \geq 1$, $u_1^2(\pi_{c_2}) = 0$ otherwise; and $w_1(u_1) = u_1$ if $u_1 \leq 1$, $w_1(u_1) = 1 + \frac{1}{100}(u_1 - 1)$ otherwise. The second bidder bids $u_2^1(\pi_{c_1}) = 1$ if $\pi_{c_1} \geq 1$, $u_1^1(\pi_{c_1}) = 0$ otherwise; $u_2^2(\pi_{c_2}) = 0$ (always); $w_2(u_2) = \frac{1}{4}u_2$ if $u_2 \leq 1$, $w_2(u_2) = \frac{1}{4} + \frac{1}{100}(u_2 - 1)$ otherwise. Now, the third bidder's *true* preferences are accurately represented[7] by the bid $u_3^1(\pi_{c_1}) = 1$ if $\pi_{c_1} \geq 1$, $u_3^1(\pi_{c_1}) = 0$ otherwise; $u_3^2(\pi_{c_2}) = 3$ if $\pi_{c_2} \geq 1$, $u_3^2(\pi_{c_1}) = 0$ otherwise; and $w_3(u_3) = \frac{1}{3}u_3$ if $u_3 \leq 1$, $w_3(u_3) = \frac{1}{3} + \frac{1}{100}(u_3 - 1)$ otherwise. Now, it is straightforward to check that, if the third bidder bids truthfully, regardless of whether the objective is surplus maximization or total donated, charity 1 will receive at least 1, and charity 2 will receive less than 1. The same is true if bidder 3 does not place a bid at all (as in the previous type of manipulation); hence bidder 2's utility will be 1 in this case. But now, if bidder 3 reports $u_3^1(\pi_{c_1}) = 0$ everywhere; $u_3^2(\pi_{c_2}) = 3$ if $\pi_{c_2} \geq 1$, $u_3^2(\pi_{c_2}) = 0$ otherwise (this part of the bid is truthful); and $w_3(u_3) = \frac{1}{3}u_3$ if $u_3 \leq 1$, $w_3(u_3) = \frac{1}{3}$ otherwise; then charity 2 will receive at least 1, and bidder 3 will have to pay at most $\frac{1}{3}$. Because up to this amount of payment, one unit of money corresponds to three units of utility to bidder 3, it follows his utility is now at least $3 - 1 = 2 > 1$. We observe that in this case, the strategic bidder is not only affecting how much the bidders pay, but also how much the charities receive.

5.2.2 Mechanism design in the quasilinear setting

In the remainder of this section, we restrict our attention to bidders with quasilinear preferences. There are at least four reasons why the mechanism design approach is likely to be most successful in the setting of quasilinear preferences. First, historically, mechanism design has been been most successful when the quasilinear assumption could be made. Second, because of this success, some very general mechanisms have been discovered for the quasilinear setting (for instance, the VCG and dAGVA mechanisms) which we could apply directly to the expressive charity donation problem (although they are not fully satisfactory, as VCG is not budget-balanced, and dAGVA is not individually rational). Third, as we saw in Section 3.3.4, the clearing problem is much easier in this setting, and thus we are less likely to run into computational trouble for the mechanism design problem. Fourth, as we will show shortly, the quasilinearity assumption in some cases allows for decomposing the mechanism design problem over the charities (as it did for the simple clearing problem).

Moreover, in the quasilinear setting (unlike in the general setting), it makes sense to pursue social welfare (the sum of the utilities) as the objective, because now 1) units of utility correspond directly to units of money, so that we do not have the problem of the bidders arbitrarily scaling their utilities; and 2) it is no longer possible to give a payment willingness function of 0 while still affecting the donations through a utility function.

We are now ready to present the result that shows that we can sometimes decompose the problem over the charities.

Theorem 35 *Suppose all agents' preferences are quasilinear (and, as we have been assuming throughout, that the utility that an agent derives from one charity is independent of how much other charities receive). Furthermore, suppose that there exists a single-charity mechanism M that, for a certain subclass P of (quasilinear) preferences, under a given solution concept S (either implementation in dominant strategies or Bayes-Nash equilibrium) and a given notion of individual*

[7] Formally, this means that if the bidder is forced to pay the full amount that his bid allows for a particular vector of payments to charities, the bidder is indifferent between this and not participating in the mechanism at all. (Compare this to bidding truthfully in a first-price auction.)

rationality R (either ex post, ex interim, or none), satisfies a certain notion of budget balance (either ex post, ex ante, or none), and is ex-post efficient. Then, there exists a mechanism with the same properties for any number of charities—namely, the mechanism that runs the single-charity mechanism separately for each individual charity.

Proof: As stated in the theorem, the mechanism is simply the following: for each charity, run the single-charity mechanism on the agents' preferences for that charity, and let the agents make the corresponding payments to that charity. (So, each agent's total payment will be the sum of her payments to the individual charities.) Because the agents are assumed to be maximizing expected utility, and the utilities that they derive from different charities are independent, it follows by linearity of expectation that they can separate their truthfulness and participation decisions across the charities. Thus, the desired properties follow from the fact that the single-charity mechanism has these properties. ■

Two mechanisms that satisfy efficiency (and can in fact be applied directly to the multiple-charity problem without use of the previous theorem) are the VCG (which is incentive compatible in dominant strategies) and dAGVA (which is incentive compatible only in Bayes-Nash equilibrium) mechanisms. Each of them, however, has a drawback that would probably make it impractical in the setting of donations to charities. The VCG mechanism is not budget balanced. The dAGVA mechanism does not satisfy ex-post individual rationality. In the next subsection, we will investigate whether we can do better in the setting of donations to charities.

5.2.3 Impossibility of efficiency

In this subsection, we show that even in a very restricted setting, and with minimal requirements on incentive-compatibility and individual-rationality constraints, it is impossible to create a mechanism that is efficient.

Theorem 36 *There is no mechanism which is ex-post budget balanced, ex-post efficient, and ex-interim individually rational with Bayes-Nash equilibrium as the solution concept (even with only one charity, only two quasilinear bidders, with identical type distributions (uniform over two types, with either both utility functions being step functions or both utility functions being concave piece-wise linear functions)).*

Proof: Suppose the two bidders both have the following distribution over types. With probability $\frac{1}{2}$, the bidder does not care for the charity at all (u is zero everywhere); otherwise, the bidder derives utility $\frac{5}{4}$ from the charity getting at least 1, and utility 0 otherwise. (Alternatively, for the second type, the bidder can get $\min\{\frac{5}{4}, \frac{5\pi_c}{4}\}$—a concave piecewise linear function.) Call the first type the low type (L), the second one the high type (H).

Suppose a mechanism with the desired properties does exist. By the revelation principle, we can assume that revealing preferences truthfully is a Bayes-Nash equilibrium in this mechanism. Because the mechanism is ex-post efficient, the charity should receive exactly 1 when either bidder has the high type, and 0 otherwise. Let $\pi_1(\theta_1, \theta_2)$ be bidder 1's (expected) payment when she reports θ_1 and the other bidder reports θ_2. By ex-interim IR, $\pi_1(L, H) + \pi_1(L, L) \leq 0$. Because bidder one cannot have an incentive to report falsely when her true type is high, we have $\frac{5}{4} - \pi_1(L,H) - \pi_1(L,L) \leq \frac{5}{4} - \pi_1(H,H) + \frac{5}{4} - \pi_1(H,L)$, or equivalently $\pi_1(H,H) + \pi_1(H,L) \leq \frac{5}{4} + \pi_1(L,L) + \pi_1(L,H) \leq \frac{5}{4}$. Because the example is completely symmetric between bidders, we can similarly conclude for bidder 2's payments that $\pi_2(H,H) + \pi_2(L,H) \leq \frac{5}{4}$. Of course, in order to pay the charity the necessary amount of 1 whenever one of the bidders has her high type, we need to have $\pi_1(H,H) + \pi_1(H,L) + \pi_2(H,H) + \pi_2(L,H) + \pi_1(L,H) + \pi_2(H,L) = 3$, and thus we can conclude that $\pi_1(L,H) + \pi_2(H,L) \geq 3 - \frac{10}{4} = \frac{1}{2}$. Because the charity receives 0 when both report low, $\pi_1(L,L) + \pi_2(L,L) = 0$ and thus we can conclude that $\pi_1(L,H) + \pi_1(L,L) + \pi_2(H,L) + \pi_2(L,L) \geq \frac{1}{2}$. But by the individual rationality constraints, $\pi_1(L,H) + \pi_1(L,L) \leq 0$ and $\pi_2(H,L) + \pi_2(L,L) \leq 0$. (Contradiction.)[8] ∎

The case of step-functions in this theorem corresponds exactly to the case of a single, fixed-size, nonexcludable public good (the "public good" being that the charity receives the desired amount)—for which such an impossibility result is already known [Mas-Colell *et al.*, 1995]. Many similar results are known, probably the most famous of which is the Myerson-Satterthwaite impossibility result, which proves the impossibility of efficient bilateral trade under the same requirements [Myerson and Satterthwaite, 1983].

Theorem 35 indicates that there is no reason to decide on donations to multiple charities under a single mechanism (rather than a separate one for each charity), when an efficient mechanism with the desired properties exists for the single-charity case. However, because under the requirements of Theorem 36, no such mechanism exists, there may be a benefit to bringing the charities under the same umbrella. The next proposition shows that this is indeed the case.

Proposition 8 *There exist settings with two charities where there exists no ex-post budget balanced, ex-post efficient, and ex-interim individually rational mechanism with Bayes-Nash equilibrium as the solution concept for either charity alone; but there exists an ex-post budget balanced, ex-post efficient, and ex-post individually rational mechanism with dominant strategies as the solution concept for both charities together. (Even when the conditions are the same as in Theorem 36, apart from the fact that there are now two charities.)*

Proof: Suppose that each bidder has two types, With probability $\frac{1}{2}$ each: for the first type, her preferences for the first charity correspond to the high type in the proof of Theorem 36, and her preferences for the second charity correspond to the low type in the proof of Theorem 36. For the second type, her preferences for the first charity correspond to the low type, and her preferences for the second charity correspond to the high type. Now, if we wish to create a mechanism for either charity individually, we are in exactly the same setting as in the proof of Theorem 36, where

[8]As an alternative proof technique (a proof by computer), we let our automated mechanism design software (described in Chapter 6) create a mechanism for the (step-function) instance described in the proof, which was restricted to be implementable in dominant strategies, ex-interim individually rational, and (weak) budget balanced, with social welfare (counting the payments made) as the objective. The mechanism did not burn any money (did not pay unnecessarily much to the charity), but did not always give money to the charity when it was beneficial to do so. (It randomized uniformly between giving 1 and giving 0 when player one's type was low, and player 2's high.) Since an ex-post budget balanced, ex-post efficient mechanism would have had a higher expected objective value, and automated mechanism design always finds the mechanism that maximizes the expected objective value under the constraints it is given, we can conclude that no ex-post budget balanced, ex-post efficient mechanism exists under the given constraints.

we know that it is impossible to get all of ex-post budget balance, ex-post efficiency, and ex-interim individually rationality in Bayes-Nash equilibrium. On the other hand, consider the following mechanism for the joint problem. If both bidders report preferring the same charity, each bidder pays $\frac{1}{2}$, and the preferred charity receives 1 (the other 0). Otherwise, each bidder pays 1, and each charity receives 1. It is straightforward to check that the mechanism is ex-post budget balanced, ex-post efficient, and ex-post individually rational. To see that truthtelling is a dominant strategy, we need to check two cases. First, if one bidder reports a high type for the charity that the other bidder does not prefer, this latter bidder is better off reporting truthfully: reporting falsely will give her utility $-\frac{1}{2}$ (nothing will be donated to her preferred utility), which is less than reporting truthfully by ex-post IR. Second, if one bidder reports a high type for the charity that the other bidder prefers, this latter bidder is better off reporting truthfully as well: her preferred charity will receive the same amount regardless of her report, but her required payment is only $\frac{1}{2}$ if she reports truthfully, as opposed to 1 if she reports falsely. ■

This concludes the part of this dissertation studying expressive preference aggregation for donations to charities.

5.3 Summary

In this chapter, we studied problems that classical mechanism design faces in some expressive preference aggregation settings. In Section 5.1, we studied two related problems concerning the VCG mechanism: the problem of revenue guarantees, and that of collusion. We studied four settings: combinatorial forward auctions with free disposal, combinatorial reverse auctions with free disposal, combinatorial forward (or reverse) auctions without free disposal, and combinatorial exchanges. In each setting, we gave an example of how additional bidders (colluders) can make the outcome much worse (less revenue or higher cost) under the VCG mechanism (but not under a first price mechanism); derived necessary and sufficient conditions for such an effective collusion to be possible under the VCG mechanism; and (when nontrivial) studied the computational complexity of deciding whether these conditions hold.

In Section 5.2, we studied mechanism design for expressive preference aggregation for donations to (charitable) causes. We showed that even with only a single charity, a fundamental impossibility result similar to the Myerson-Satterthwaite impossibility theorem holds; but we also showsed some positive results, including how mechanisms that are successful in single-charity settings can be extended to settings with multiple charities, and how combining the aggregation of preferences over donations to multiple individual charities into a single mechanism can improve efficiency.

The work in this chapter provides some reasons why simply taking a standard mechanism "off the shelf" is not always satisfactory, especially in domains with complex preferences. Rather, it may be preferable to design a custom mechanism. The next chapter takes this idea to its extreme: we will study how an optimal mechanism can be automatically designed (computed) *for the specific instance at hand only*.