

Expressive Negotiation over Donations to Charities*

Vincent Conitzer
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
conitzer@cs.cmu.edu

Tuomas Sandholm
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
sandholm@cs.cmu.edu

ABSTRACT

When donating money to a (say, charitable) cause, it is possible to use the contemplated donation as *negotiating material* to induce other parties interested in the charity to donate more. Such negotiation is usually done in terms of *matching offers*, where one party promises to pay a certain amount if others pay a certain amount. However, in their current form, matching offers allow for only limited negotiation. For one, it is not immediately clear how multiple parties can make matching offers at the same time without creating circular dependencies. Also, it is not immediately clear how to make a donation conditional on other donations to multiple charities, when the donator has different levels of appreciation for the different charities. In both these cases, the limited expressiveness of matching offers causes economic loss: it may happen that an arrangement that would have made all parties (donators as well as charities) better off cannot be expressed in terms of matching offers and will therefore not occur.

In this paper, we introduce a bidding language for expressing very general types of matching offers over multiple charities. We formulate the corresponding clearing problem (deciding how much each bidder pays, and how much each charity receives), and show that it is NP-complete to approximate to any ratio even in very restricted settings. We give a mixed-integer program formulation of the clearing problem, and show that for concave bids, the program reduces to a linear program. We then show that the clearing problem for a subclass of concave bids is at least as hard as the decision variant of linear programming. Subsequently, we show that the clearing problem is much easier when bids are quasilinear—for surplus, the problem decomposes across charities, and for payment maximization, a greedy approach is optimal if the bids are concave (although this latter problem is weakly NP-complete when the bids are not concave). For the quasilinear setting, we study the mechanism design question. We show that an ex-post efficient mechanism is

impossible even with only one charity and a very restricted class of bids. We also show that there may be benefits to linking the charities from a mechanism design standpoint.

Categories and Subject Descriptors

F.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithms, Economics, Theory

Keywords

Expressive Negotiation, Donations to Charities, Market Clearing, Mechanism Design

1. INTRODUCTION

When money is donated to a charitable (or other) cause (hereafter referred to as *charity*), often the donating party gives *unconditionally*: a fixed amount is transferred from the donator to the charity, and none of this transfer is contingent on other events—in particular, it is not contingent on the amount given by other parties. Indeed, this is currently often the only way to make a donation, especially for small donating parties such as private individuals. However, when multiple parties support the same charity, each of them would prefer to see the others give more rather than less to this charity. In such scenarios, it is sensible for a party to use its contemplated donation as *negotiating material* to induce the others to give more. This is done by making the donation conditional on the others' donations. The following example will illustrate this, and show that the donating parties as well as the charitable cause may simultaneously benefit from the potential for such negotiation.

Suppose we have two parties, 1 and 2, who are both supporters of charity *A*. To either of them, it would be worth \$0.75 if *A* received \$1. It follows neither of them will be willing to give unconditionally, because $\$0.75 < \1 . However, if the two parties draw up a contract that says that they will each give \$0.5, both the parties have an incentive to accept this contract (rather than have no contract at all): with the contract, the charity will receive \$1 (rather than \$0 without a contract), which is worth \$0.75 to each party, which is greater than the \$0.5 that that party will have to give. Effectively, each party has made its donation conditional on the other party's donation, leading to larger donations and greater happiness to all parties involved.

*Supported by NSF under CAREER Award IRI-9703122, Grant IIS-9800994, ITR IIS-0081246, and ITR IIS-0121678.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'04, May 17–20, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-711-0/04/0005 ...\$5.00.

One method that is often used to effect this is to make a *matching offer*. Examples of matching offers are: “I will give x dollars for every dollar donated.”, or “I will give x dollars if the total collected from other parties exceeds y .” In our example above, one of the parties can make the offer “I will donate \$0.5 if the other party also donates at least that much”, and the other party will have an incentive to indeed donate \$0.5, so that the total amount given to the charity increases by \$1. Thus this matching offer implements the contract suggested above. As a real-world example, the United States government has authorized a donation of up to \$1 billion to the Global Fund to fight AIDS, TB and Malaria, under the condition that the American contribution does not exceed one third of the total—to encourage other countries to give more [23].

However, there are several severe limitations to the simple approach of matching offers as just described.

1. It is not clear how two parties can make matching offers where each party’s offer is stated in terms of the amount that the other pays. (For example, it is not clear what the outcome should be when both parties offer to match the other’s donation.) Thus, matching offers can only be based on payments made by parties that are giving unconditionally (not in terms of a matching offer)—or at least there can be no circular dependencies.¹
2. Given the current infrastructure for making matching offers, it is impractical to make a matching offer depend on the amounts given to *multiple* charities. For instance, a party may wish to specify that it will pay \$100 given that charity A receives a total of \$1000, but that it will also count donations made to charity B , at half the rate. (Thus, a total payment of \$500 to charity A combined with a total payment of \$1000 to charity B would be just enough for the party’s offer to take effect.)

In contrast, in this paper we propose a new approach where each party can express its relative preferences for different charities, and make its offer conditional on its own appreciation for the vector of donations made to the different charities. Moreover, the amount the party offers to donate at different levels of appreciation is allowed to vary arbitrarily (it does need to be a dollar-for-dollar (or n -dollar-for-dollar) matching arrangement, or an arrangement where the party offers a fixed amount provided a given (strike) total has been exceeded). Finally, there is a clear interpretation of what it means when multiple parties are making conditional offers that are stated in terms of each other. Given each combination of (conditional) offers, there is a (usually) unique solution which determines how much each party pays, and how much each charity is paid.

However, as we will show, finding this solution (the *clearing problem*) requires solving a potentially difficult optimization problem. A large part of this paper is devoted to studying how difficult this problem is under different assumptions on the structure of the offers, and providing algorithms for solving it.

¹Typically, larger organizations match offers of private individuals. For example, the American Red Cross Liberty Disaster Fund maintains a list of businesses that match their customers’ donations [8].

Towards the end of the paper, we also study the *mechanism design* problem of motivating the bidders to bid truthfully.

In short, expressive negotiation over donations to charities is a new way in which electronic commerce can help the world. A web-based implementation of the ideas described in this paper can facilitate voluntary reallocation of wealth on a global scale. Additionally, optimally solving the clearing problem (and thereby generating the maximum economic welfare) requires the application of sophisticated algorithms.

2. COMPARISON TO COMBINATORIAL AUCTIONS AND EXCHANGES

This section discusses the relationship between expressive charity donation and combinatorial auctions and exchanges. It can be skipped, but may be of interest to the reader with a background in combinatorial auctions and exchanges.

In a *combinatorial auction*, there are m items for sale, and bidders can place bids on *bundles* of one or more items. The auctioneer subsequently labels each bid as winning or losing, under the constraint that no item can be in more than one winning bid, to maximize the sum of the values of the winning bids. (This is known as the *clearing problem*.) Variants include combinatorial *reverse* auctions, where the auctioneer is seeking to procure a set of items; and combinatorial *exchanges*, where bidders can both buy and sell items (even within the same bid). Other extensions include allowing for *side constraints*, as well as the specification of *attributes* of the items in bids. Combinatorial auctions and exchanges have recently become a popular research topic [20, 21, 17, 22, 9, 18, 13, 3, 12, 26, 19, 25, 2].

The problems of clearing expressive charity donation markets and clearing combinatorial auctions or exchanges are very different in formulation. Nevertheless, there are interesting parallels. One of the main reasons for the interest in combinatorial auctions and exchanges is that it allows for *expressive bidding*. A bidder can express exactly how much each different *allocation* is worth to her, and thus the globally optimal allocation may be chosen by the auctioneer. Compare this to a bidder having to bid on two different items in two different (one-item) auctions, without any way of expressing that (for instance) one item is worthless if the other item is not won. In this scenario, the bidder may win the first item but not the second (because there was another high bid on the second item that she did not anticipate), leading to economic inefficiency.

Expressive bidding is also one of the main benefits of the expressive charity donation market. Here, bidders can express exactly how much they are willing to donate for every vector of amounts donated to charities. This may allow bidders to negotiate a complex arrangement of who gives how much to which charity, which is beneficial to all parties involved; whereas no such arrangement may have been possible if the bidders had been restricted to using simple matching offers on individual charities. Again, expressive bidding is necessary to achieve economic efficiency.

Another parallel is the computational complexity of the clearing problem. In order to achieve the full economic efficiency allowed by the market’s expressiveness (or even come close to it), hard computational problems must be solved in combinatorial auctions and exchanges, as well as in the charity donation market (as we will see).

3. DEFINITIONS

Throughout this paper, we will refer to the offers that the donating parties make as *bids*, and to the donating parties as *bidders*. In our bidding framework, a bid will specify, for each vector of total payments made to the charities, how much that bidder is willing to contribute. (The contribution of this bidder is also counted in the vector of payments—so, the vector of total payments to the charities represents the amount given by *all* donating parties, not just the ones other than this bidder.) The bidding language is expressive enough that no bidder should have to make more than one bid. The following definition makes the general form of a bid in our framework precise.

DEFINITION 1. *In a setting with m charities c_1, c_2, \dots, c_m , a bid by bidder b_j is a function $v_j : \mathbb{R}^m \rightarrow \mathbb{R}$. The interpretation is that if charity c_i receives a total amount of π_{c_i} , then bidder j is willing to donate (up to) $v_j(\pi_{c_1}, \pi_{c_2}, \dots, \pi_{c_m})$.*

We now define possible outcomes in our model, and which outcomes are valid given the bids that were made.

DEFINITION 2. *An outcome is a vector of payments made by the bidders $(\pi_{b_1}, \pi_{b_2}, \dots, \pi_{b_n})$, and a vector of payments received by the charities $(\pi_{c_1}, \pi_{c_2}, \dots, \pi_{c_m})$. A valid outcome is an outcome where*

1. $\sum_{j=1}^n \pi_{b_j} \geq \sum_{i=1}^m \pi_{c_i}$ (at least as much money is collected as is given away);
2. For all $1 \leq j \leq n$, $\pi_{b_j} \leq v_j(\pi_{c_1}, \pi_{c_2}, \dots, \pi_{c_m})$ (no bidder gives more than she is willing to).

Of course, in the end, only one of the valid outcomes can be chosen. We choose the valid outcome that maximizes the objective that we have for the donation process.

DEFINITION 3. *An objective is a function from the set of all outcomes to \mathbb{R} .² After all bids have been collected, a valid outcome will be chosen that maximizes this objective.*

One example of an objective is *surplus*, given by $\sum_{j=1}^n \pi_{b_j} - \sum_{i=1}^m \pi_{c_i}$. The surplus could be the profits of a company managing the expressive donation marketplace; but, alternatively, the surplus could be returned to the bidders, or given to the charities. Another objective is *total amount donated*, given by $\sum_{i=1}^m \pi_{c_i}$. (Here, different weights could also be placed on the different charities.)

Finding the valid outcome that maximizes the objective is a (nontrivial) computational problem. We will refer to it as the *clearing* problem. The formal definition follows.

DEFINITION 4 (DONATION-CLEARING). *We are given a set of n bids over charities c_1, c_2, \dots, c_m . Additionally, we are given an objective function. We are asked to find an objective-maximizing valid outcome.*

How difficult the DONATION-CLEARING problem is depends on the types of bids used and the language in which they are expressed. This is the topic of the next section.

²In general, the objective function may also depend on the bids, but the objective functions under consideration in this paper do not depend on the bids. The techniques presented in this paper will typically generalize to objectives that take the bids into account directly.

4. A SIMPLIFIED BIDDING LANGUAGE

Specifying a general bid in our framework (as defined above) requires being able to specify an arbitrary real-valued function over \mathbb{R}^m . Even if we restricted the possible total payment made to each charity to the set $\{0, 1, 2, \dots, s\}$, this would still require a bidder to specify $(s+1)^m$ values. Thus, we need a bidding language that will allow the bidders to at least specify *some* bids more concisely. We will specify a bidding language that only represents a subset of all possible bids, which can be described concisely.³

To introduce our bidding language, we will first describe the bidding function as a composition of two functions; then we will outline our assumptions on each of these functions. First, there is a *utility function* $u_j : \mathbb{R}^m \rightarrow \mathbb{R}$, specifying how much bidder j appreciates a given vector of total donations to the charities. (Note that the way we define a bidder's utility function, it does not take the payments the bidder makes into account.) Then, there is a *donation willingness* function $w_j : \mathbb{R} \rightarrow \mathbb{R}$, which specifies how much bidder j is willing to pay given her utility for the vector of donations to the charities. We emphasize that this function does *not* need to be linear, so that utilities should not be thought of as expressible in dollar amounts. (Indeed, when an individual is donating to a large charity, the reason that the individual donates only a bounded amount is typically not decreasing marginal value of the money given to the charity, but rather that the marginal value of a dollar to the bidder herself becomes larger as her budget becomes smaller.) So, we have $w_j(u_j(\pi_{c_1}, \pi_{c_2}, \dots, \pi_{c_m})) = v_j(\pi_{c_1}, \pi_{c_2}, \dots, \pi_{c_m})$, and we let the bidder describe her functions u_j and w_j separately. (She will submit these functions as her bid.)

Our first restriction is that the utility that a bidder derives from money donated to one charity is *independent* of the amount donated to another charity. Thus,

$u_j(\pi_{c_1}, \pi_{c_2}, \dots, \pi_{c_m}) = \sum_{i=1}^m u_j^i(\pi_{c_i})$. (We observe that this does *not* imply that the bid function v_j decomposes similarly, because of the nonlinearity of w_j .) Furthermore, each u_j^i must be piecewise linear. An interesting special case which we will study is when each u_j^i is a line: $u_j^i(\pi_{c_i}) = a_j^i \pi_{c_i}$. This special case is justified in settings where the scale of the donations by the bidders is small relative to the amounts the charities receive from other sources, so that the marginal use of a dollar to the charity is not affected by the amount given by the bidders.

The only restriction that we place on the payment willingness functions w_j is that they are piecewise linear. One interesting special case is a *threshold bid*, where w_j is a step function: the bidder will provide t dollars if her utility exceeds s , and otherwise 0. Another interesting case is when such a bid is *partially acceptable*: the bidder will provide t dollars if her utility exceeds s ; but if her utility is $u < s$, she is still willing to provide $\frac{u-s}{s}$ dollars.

One might wonder why, if we are given the bidders' utility functions, we do not simply maximize the sum of the utilities rather than surplus or total donated. There are several reasons. First, because affine transformations do not affect utility functions in a fundamental way, it would be possi-

³Of course, our bidding language can be trivially extended to allow for fully expressive bids, by also allowing bids from a fully expressive bidding language, in addition to the bids in our bidding language.

ble for a bidder to inflate her utility by changing its units, thereby making her bid more important for utility maximization purposes. Second, a bidder could simply give a payment willingness function that is 0 everywhere, and have her utility be taken into account in deciding on the outcome, in spite of her not contributing anything.

5. AVOIDING INDIRECT PAYMENTS

In an initial implementation, the approach of having donations made out to a center, and having a center forward these payments to charities, may not be desirable. Rather, it may be preferable to have a *partially decentralized* solution, where the donating parties write out checks to the charities directly according to a solution prescribed by the center. In this scenario, the center merely has to verify that parties are giving the prescribed amounts. Advantages of this include that the center can keep its legal status minimal, as well as that we do not require the donating parties to trust the center to transfer their donations to the charities (or require some complicated verification protocol). It is also a step towards a fully decentralized solution, if this is desirable.

To bring this about, we can still use the approach described earlier. After we clear the market in the manner described before, we know the amount that each donor is supposed to give, and the amount that each charity is supposed to receive. Then, it is straightforward to give some specification of who should give how much to which charity, that is consistent with that clearing. Any greedy algorithm that increases the cash flow from any bidder who has not yet paid enough, to any charity that has not yet received enough, until either the bidder has paid enough or the charity has received enough, will provide such a specification. (All of this is assuming that $\sum_{b_j} \pi_{b_j} = \sum_{c_i} \pi_{c_i}$. In the case

where there is nonzero surplus, that is, $\sum_{b_j} \pi_{b_j} > \sum_{c_i} \pi_{c_i}$, we can distribute this surplus across the bidders by not requiring them to pay the full amount, or across the charities by giving them more than the solution specifies.)

Nevertheless, with this approach, a bidder may have to write out a check to a charity that she does not care for at all. (For example, an environmental activist who was using the system to increase donations to a wildlife preservation fund may be required to write a check to a group supporting a right-wing political party.) This is likely to lead to complaints and noncompliance with the clearing. We can address this issue by letting each bidder specify explicitly (before the clearing) which charities she would be willing to make a check out to. These additional constraints, of course, may change the optimal solution. In general, checking whether a given centralized solution (with zero surplus) can be accomplished through decentralized payments when there are such constraints can be modeled as a MAX-FLOW problem. In the MAX-FLOW instance, there is an edge from the source node s to each bidder b_j , with a capacity of π_{b_j} (as specified in the centralized solution); an edge from each bidder b_j to each charity c_i that the bidder is willing to donate money to, with a capacity of ∞ ; and an edge from each charity c_i to the target node t with capacity π_{c_i} (as specified in the centralized solution).

In the remainder of this paper, all our hardness results apply even to the setting where there is no constraint on which bidders can pay to which charity (that is, even the problem

as it was specified before this section is hard). We also generalize our clearing algorithms to the partially decentralized case with constraints.

6. HARDNESS OF CLEARING THE MARKET

In this section, we will show that the clearing problem is completely inapproximable, even when every bidder's utility function is linear (with slope 0 or 1 in each charity's payments), each bidder cares either about at most two charities or about all charities equally, and each bidder's payment willingness function is a step function. We will reduce from MAX2SAT (given a formula in conjunctive normal form (where each clause has two literals) and a target number of satisfied clauses T , does there exist an assignment of truth values to the variables that makes at least T clauses true?), which is NP-complete [7].

THEOREM 1. *There exists a reduction from MAX2SAT instances to DONATION-CLEARING instances such that*

1. *If the MAX2SAT instance has no solution, then the only valid outcome is the zero outcome (no bidder pays anything and no charity receives anything);*
2. *Otherwise, there exists a solution with positive surplus. Additionally, the DONATION-CLEARING instances that we reduce to have the following properties:*
 1. *Every u_j^i is a line; that is, the utility that each bidder derives from any charity is linear;*
 2. *All the u_j^i have slope either 0 or 1;*
 3. *Every bidder either has at most 2 charities that affect her utility (with slope 1), or all charities affect her utility (with slope 1);*
 4. *Every bid is a threshold bid; that is, every bidder's payment willingness function w_j is a step function.*

PROOF. The problem is in NP because we can nondeterministically choose the payments to be made and received, and check the validity and objective value of this outcome.

In the following, we will represent bids as follows: $(\{(c_k, a_k)\}, s, t)$ indicates that $u_j^k(\pi_{c_k}) = a_k \pi_{c_k}$ (this function is 0 for c_k not mentioned in the bid), and $w_j(u_j) = t$ for $u_j \geq s$, $w_j(u_j) = 0$ otherwise.

To show NP-hardness, we reduce an arbitrary MAX2SAT instance, given by a set of clauses $K = \{k\} = \{(l_k^1, l_k^2)\}$ over a variable set V together with a target number of satisfied clauses T , to the following DONATION-CLEARING instance. Let the set of charities be as follows. For every literal $l \in L$, there is a charity c_l . Then, let the set of bids be as follows. For every variable v , there is a bid $b_v = (\{(c_{+v}, 1), (c_{-v}, 1)\}, 2, 1 - \frac{1}{4|V|})$. For every literal l , there is a bid $b_l = (\{(c_l, 1)\}, 2, 1)$. For every clause $k = \{l_k^1, l_k^2\} \in K$, there is a bid $b_k = (\{(c_{l_k^1}, 1), (c_{l_k^2}, 1)\}, 2, \frac{1}{8|V||K|})$. Finally, there is a single bid that values all charities equally: $b_0 = (\{(c_1, 1), (c_2, 1), \dots, (c_m, 1)\}, 2|V| + \frac{T}{8|V||K|}, \frac{1}{4} + \frac{1}{16|V||K|})$. We show the two instances are equivalent.

First, suppose there exists a solution to the MAX2SAT instance. If in this solution, l is *true*, then let $\pi_{c_l} = 2 + \frac{T}{8|V||K|}$; otherwise $\pi_{c_l} = 0$. Also, the only bids that are *not* accepted (meaning the threshold is not met) are the b_l where l is *false*, and the b_k such that both of l_k^1, l_k^2 are *false*. First we show that no bidder whose bid is accepted pays more than she is willing to. For each b_v , either c_{+v} or c_{-v} receives at least 2, so this bidder's threshold has been met.

For each b_l , either l is *false* and the bid is not accepted, or l is *true*, c_l receives at least 2, and the threshold has been met. For each b_k , either both of l_k^1, l_k^2 are *false* and the bid is not accepted, or at least one of them (say l_k^1) is *true* (that is, k is satisfied) and $c_{l_k^1}$ receives at least 2, and the threshold has been met. Finally, because the total amount received by the charities is $2|V| + \frac{T}{8|V||K|}$, b_0 's threshold has also been met. The total amount that can be extracted from the accepted bids is at least $|V|(1 - \frac{1}{4|V|}) + |V| + T \frac{1}{8|V||K|} + \frac{1}{4} + \frac{1}{16|V||K|} = 2|V| + \frac{T}{8|V||K|} + \frac{1}{16|V||K|} > 2|V| + \frac{T}{8|V||K|}$, so there is positive surplus. So there exists a solution with positive surplus to the DONATION-CLEARING instance.

Now suppose there exists a nonzero outcome in the DONATION-CLEARING instance. First we show that it is not possible (for any $v \in V$) that both b_{+v} and b_{-v} are accepted. For, this would require that $\pi_{c_{+v}} + \pi_{c_{-v}} \geq 4$. The bids b_v, b_{+v}, b_{-v} cannot contribute more than 3, so we need another 1 at least. It is easily seen that for any other v' , accepting any subset of $\{b_{v'}, b_{+v'}, b_{-v'}\}$ would require that at least as much is given to $c_{+v'}$ and $c_{-v'}$ as can be extracted from these bids, so this cannot help. Finally, all the other bids combined can contribute at most $|K| \frac{1}{8|V||K|} + \frac{1}{4} + \frac{1}{16|V||K|} < 1$. It follows that we can interpret the outcome in the DONATION-CLEARING instance as a partial assignment of truth values to variables: v is set to *true* if b_{+v} is accepted, and to *false* if b_{-v} is accepted. All that is left to show is that this partial assignment satisfies at least T clauses.

First we show that if a clause bid b_k is accepted, then either $b_{l_k^1}$ or $b_{l_k^2}$ is accepted (and thus either l_k^1 or l_k^2 is set to *true*, hence k is satisfied). If b_k is accepted, at least one of $c_{l_k^1}$ and $c_{l_k^2}$ must be receiving at least 1; without loss of generality, say it is $c_{l_k^1}$, and say l_k^1 corresponds to variable v_k^1 (that is, it is $+v_k^1$ or $-v_k^1$). If $c_{l_k^1}$ does not receive at least 2, $b_{l_k^1}$ is not accepted, and it is easy to check that the bids $b_{v_k^1}, b_{+v_k^1}, b_{-v_k^1}$ contribute (at least) 1 less than is paid to $c_{+v_k^1}$ and $c_{-v_k^1}$. But this is the same situation that we analyzed before, and we know it is impossible. All that remains to show is that at least T clause bids are accepted.

We now show that b_0 is accepted. Suppose it is not; then one of the b_v must be accepted. (The solution is nonzero by assumption; if only some b_k are accepted, the total payment from these bids is at most $|K| \frac{1}{8|V||K|} < 1$, which is not enough for any bid to be accepted; and if one of the b_l is accepted, then the threshold for the corresponding b_v is also reached.) For this v , $b_{v^1}, b_{+v^1}, b_{-v^1}$ contribute (at least) $\frac{1}{4|V|}$ less than the total payments to c_{+v} and c_{-v} . Again, the other b_v and b_l cannot (by themselves) help to close this gap; and the b_k can contribute at most $|K| \frac{1}{8|V||K|} < \frac{1}{4|V|}$. It follows that b_0 is accepted.

Now, in order for b_0 to be accepted, a total of $2|V| + \frac{T}{8|V||K|}$ must be donated. Because it is not possible (for any $v \in V$) that both b_{+v} and b_{-v} are accepted, it follows that the total payment by the b_v and the b_l can be at most $2|V| - \frac{1}{4}$. Adding b_0 's payment of $\frac{1}{4} + \frac{1}{16|V||K|}$ to this, we still need $\frac{T - \frac{1}{2}}{8|V||K|}$ from the b_k . But each one of them contributes at most $\frac{1}{8|V||K|}$, so at least T of them must be accepted. \square

COROLLARY 1. *Unless $P=NP$, there is no polynomial-time algorithm for approximating DONATION-CLEARING (with either the surplus or the total amount donated as the objective) within any ratio $f(n)$, where f is a nonzero function of the size of the instance. This holds even if the DONATION-CLEARING structures satisfy all the properties given in Theorem 1.*

PROOF. Suppose we had such a polynomial time algorithm, and applied it to the DONATION-CLEARING instances that were reduced from MAX2SAT instances in Theorem 1. It would return a nonzero solution when the MAX2SAT instance has a solution, and a zero solution otherwise. So we can decide whether arbitrary MAX2SAT instances are satisfiable this way, and it would follow that $P=NP$. \square

(Solving the problem to optimality is NP-complete in many other (noncomparable or even more restricted) settings as well—we omit such results because of space constraint.) This should not be interpreted to mean that our approach is infeasible. First, as we will show, there are very expressive families of bids for which the problem is solvable in polynomial time. Second, NP-completeness is often overcome in practice (especially when the stakes are high). For instance, even though the problem of clearing combinatorial auctions is NP-complete [20] (even to approximate [21]), they are typically solved to optimality in practice.

7. MIXED INTEGER PROGRAMMING FORMULATION

In this section, we give a mixed integer programming (MIP) formulation for the general problem. We also discuss in which special cases this formulation reduces to a linear programming (LP) formulation. In such cases, the problem is solvable in polynomial time, because linear programs can be solved in polynomial time [11].

The variables of the MIP defining the final outcome are the payments made to the charities, denoted by π_{c_i} , and the payments extracted from the bidders, π_{b_j} . In the case where we try to avoid direct payments and let the bidders pay the charities directly, we add variables π_{c_i, b_j} indicating how much b_j pays to c_i , with the constraints that for each c_i , $\pi_{c_i} \leq \sum_{b_j} \pi_{c_i, b_j}$; and for each b_j , $\pi_{b_j} \geq \sum_{c_i} \pi_{c_i, b_j}$. Additionally, there is a constraint $\pi_{c_i, b_j} = 0$ whenever bidder b_j is unwilling to pay charity c_i . The rest of the MIP can be phrased in terms of the π_{c_i} and π_{b_j} .

The objectives we have discussed earlier are both linear: surplus is given by $\sum_{j=1}^n \pi_{b_j} - \sum_{i=1}^m \pi_{c_i}$, and total amount donated is given by $\sum_{i=1}^m \pi_{c_i}$ (coefficients can be added to represent different weights on the different charities in the objective).

The constraint that the outcome should be valid (no deficit) is given simply by: $\sum_{j=1}^n \pi_{b_j} \geq \sum_{i=1}^m \pi_{c_i}$.

For every bidder, for every charity, we define an additional utility variable u_j^i indicating the utility that this bidder derives from the payment to this charity. The bidder's total

utility is given by another variable u_j , with the constraint that $u_j = \sum_{i=1}^m u_j^i$.

Each u_j^i is given as a function of π_{c_i} by the (piecewise linear) function provided by the bidder. In order to represent this function in the MIP formulation, we will merely place upper bounding constraints on u_j^i , so that it cannot exceed the given functions. The MIP solver can then push the u_j^i variables all the way up to the constraint, in order to extract as much payment from this bidder as possible. In the case where the u_j^i are concave, this is easy: if (s_l, t_l) and (s_{l+1}, t_{l+1}) are endpoints of a finite linear segment in the function, we add the constraint that $u_j^i \leq t_l + \frac{\pi_{c_i} - s_l}{s_{l+1} - s_l} (t_{l+1} - t_l)$. If the final (infinite) segment starts at (s_k, t_k) and has slope d , we add the constraint that $u_j^i \leq t_k + d(\pi_{c_i} - s_k)$. Using the fact that the function is concave, for each value of π_{c_i} , the tightest upper bound on u_j^i is the one corresponding to the segment above that value of π_{c_i} , and therefore these constraints are sufficient to force the correct value of u_j^i .

When the function is not concave, we require (for the first time) some binary variables. First, we define another point on the function: $(s_{k+1}, t_{k+1}) = (s_k + M, t_k + dM)$, where d is the slope of the infinite segment and M is any upper bound on the π_{c_j} . This has the effect that we will never be on the infinite segment again. Now, let $x_l^{i,j}$ be an indicator variable that should be 1 if π_{c_i} is below the l th segment of the function, and 0 otherwise. To effect this, first add a constraint $\sum_{l=0}^k x_l^{i,j} = 1$. Now, we aim to represent π_{c_i} as a weighted average of its two neighboring $s_l^{i,j}$. For $0 \leq l \leq k+1$, let $\lambda_l^{i,j}$ be the weight on $s_l^{i,j}$. We add the constraint $\sum_{l=0}^{k+1} \lambda_l^{i,j} = 1$. Also, for $0 \leq l \leq k+1$, we add the constraint $\lambda_l^{i,j} \leq x_{l-1} + x_l$ (where x_{-1} and x_{k+1} are defined to be zero), so that indeed only the two neighboring $s_l^{i,j}$ have nonzero weight. Now we add the constraint $\pi_{c_i} = \sum_{l=0}^{k+1} s_l^{i,j} \lambda_l^{i,j}$, and now the $\lambda_l^{i,j}$ must be set correctly. Then, we can set $u_j^i = \sum_{l=0}^{k+1} t_l^{i,j} \lambda_l^{i,j}$. (This is a standard MIP technique [16].)

Finally, each π_{b_j} is bounded by a function of u_j by the (piecewise linear) function provided by the bidder (w_j). Representing this function is entirely analogous to how we represented u_j^i as a function of π_{c_i} . (Again we will need binary variables only if the function is not concave.)

Because we only use binary variables when either a utility function u_j^i or a payment willingness function w_j is not concave, it follows that if all of these are concave, our MIP formulation is simply a linear program—which can be solved in polynomial time. Thus:

THEOREM 2. *If all functions u_j^i and w_j are concave (and piecewise linear), the DONATION-CLEARING problem can be solved in polynomial time using linear programming.*

Even if some of these functions are not concave, we can simply replace each such function by the smallest upper bounding concave function, and use the linear programming formulation to obtain an upper bound on the objective—which may be useful in a search formulation of the general problem.

8. WHY ONE CANNOT DO MUCH BETTER THAN LINEAR PROGRAMMING

One may wonder if, for the special cases of the DONATION-CLEARING problem that can be solved in polynomial time with linear programming, there exist special purpose algorithms that are much faster than linear programming algorithms. In this section, we show that this is not the case. We give a reduction *from* (the decision variant of) the general linear programming problem to (the decision variant of) a special case of the DONATION-CLEARING problem (which can be solved in polynomial time using linear programming). (The decision variant of an optimization problem asks the binary question: “Can the objective value exceed o ?”) Thus, any special-purpose algorithm for solving the decision variant of this special case of the DONATION-CLEARING problem could be used to solve a decision question about an arbitrary linear program just as fast. (And thus, if we are willing to call the algorithm a logarithmic number of times, we can solve the optimization version of the linear program.)

We first observe that for linear programming, a decision question about the objective can simply be phrased as another constraint in the LP (forcing the objective to exceed the given value); then, the original decision question coincides with asking whether the resulting linear program has a feasible solution.

THEOREM 3. *The question of whether an LP (given by a set of linear constraints⁴) has a feasible solution can be modeled as a DONATION-CLEARING instance with payment maximization as the objective, with $2v$ charities and $v + c$ bids (where v is the number of variables in the LP, and c is the number of constraints). In this model, each bid b_j has only linear u_j^i functions, and is a partially acceptable threshold bid ($w_j(u) = t_j$ for $u \geq s_j$, otherwise $w_j(u) = \frac{ut_j}{s_j}$). The v bids corresponding to the variables mention only two charities each; the c bids corresponding to the constraints mention only two times the number of variables in the corresponding constraint.*

PROOF. For every variable x_i in the LP, let there be two charities, c_{+x_i} and c_{-x_i} . Let H be some number such that if there is a feasible solution to the LP, there is one in which every variable has absolute value at most H .

In the following, we will represent bids as follows: $(\{(c_k, a_k)\}, s, t)$ indicates that $u_j^k(\pi_{c_k}) = a_k \pi_{c_k}$ (this function is 0 for c_k not mentioned in the bid), and $w_j(u_j) = t$ for $u_j \geq s$, $w_j(u_j) = \frac{u_j t}{s}$ otherwise.

For every variable x_i in the LP, let there be a bid $b_{x_i} = (\{(c_{+x_i}, 1), (c_{-x_i}, 1)\}, 2H, 2H - \frac{c}{v})$. For every constraint $\sum_i r_i^j x_i \leq s_j$ in the linear program, let there be a bid $b_j = (\{(c_{-x_i}, r_i^j)\}_{i:r_i^j > 0} \cup \{(c_{+x_i}, -r_i^j)\}_{i:r_i^j < 0}, (\sum_i |r_i^j|)H - s_j, 1)$.

Let the target total amount donated be $2vH$.

Suppose there is a feasible solution $(x_1^*, x_2^*, \dots, x_v^*)$ to the LP. Without loss of generality, we can suppose that $|x_i^*| \leq H$ for all i . Then, in the DONATION-CLEARING instance,

⁴These constraints must include bounds on the variables (including nonnegativity bounds), if any.

for every i , let $\pi_{c+x_i} = H + x_i^*$, and let $\pi_{c-x_i} = H - x_i^*$ (for a total payment of $2H$ to these two charities). This allows us to extract the maximum payment from the bids b_{x_i} —a total payment of $2vH - c$. Additionally, the utility of bidder b_j is now $\sum_{i:r_i^j > 0} r_i^j(H - x_i^*) + \sum_{i:r_i^j < 0} -r_i^j(H + x_i^*) = (\sum_i |r_i^j|)H - \sum_i r_i^j x_i^* \geq (\sum_i |r_i^j|)H - s_j$ (where the last inequality stems from the fact that constraint j must be satisfied in the LP solution), so it follows we can extract the maximum payment from all the bidders b_j , for a total payment of c . It follows that we can extract the required $2vH$ payment from the bidders, and there exists a solution to the DONATION-CLEARING instance with a total amount donated of at least $2vH$.

Now suppose there is a solution to the DONATION-CLEARING instance with a total amount donated of at least vH . Then the maximum payment must be extracted from each bidder. From the fact that the maximum payment must be extracted from each bidder b_{x_i} , it follows that for each i , $\pi_{c+x_i} + \pi_{c-x_i} \geq 2H$. Because the maximum extractable total payment is $2vH$, it follows that for each i , $\pi_{c+x_i} + \pi_{c-x_i} = 2H$. Let $x_i^* = \pi_{c+x_i} - H = H - \pi_{c-x_i}$. Then, from the fact that the maximum payment must be extracted from each bidder b_j , it follows that $(\sum_i |r_i^j|)H - s_j \leq \sum_{i:r_i^j > 0} r_i^j \pi_{c-x_i} + \sum_{i:r_i^j < 0} -r_i^j \pi_{c+x_i} = \sum_{i:r_i^j > 0} r_i^j(H - x_i^*) + \sum_{i:r_i^j < 0} -r_i^j(H + x_i^*) = (\sum_i |r_i^j|)H - \sum_i r_i^j x_i^*$. Equivalently, $\sum_i r_i^j x_i^* \leq s_j$. It follows that the x_i^* constitute a feasible solution to the LP. \square

9. QUASILINEAR BIDS

Another class of bids of interest is the class of *quasilinear bids*. In a quasilinear bid, the bidder's payment willingness function is linear in utility: that is, $w_j = u_j$. (Because the units of utility are arbitrary, we may as well let them correspond exactly to units of money—so we do not need a constant multiplier.) In most cases, quasilinearity is an unreasonable assumption: for example, usually bidders have a limited budget for donations, so that the payment willingness will stop increasing in utility after some point (or at least increase slower in the case of a “softer” budget constraint). Nevertheless, quasilinearity may be a reasonable assumption in the case where the bidders are large organizations with large budgets, and the charities are a few small projects requiring relatively little money. In this setting, once a certain small amount has been donated to a charity, a bidder will derive no more utility from more money being donated from that charity. Thus, the bidders will never reach a high enough utility for their budget constraint (even when it is soft) to take effect, and thus a linear approximation of their payment willingness function is reasonable. Another reason for studying the quasilinear setting is that it is the easiest setting for mechanism design, which we will discuss shortly. In this section, we will see that the clearing problem is much easier in the case of quasilinear bids.

First, we address the case where we are trying to maximize surplus (which is the most natural setting for mechanism design). The key observation here is that when bids are quasilinear, the clearing problem *decomposes* across charities.

LEMMA 1. *Suppose all bids are quasilinear, and surplus is the objective. Then we can clear the market optimally by clearing the market for each charity individually. That is, for each bidder b_j , let $\pi_{b_j} = \sum_{c_i} \pi_{b_j^i}$. Then, for each charity c_i , maximize $(\sum_{b_j} \pi_{b_j^i}) - \pi_{c_i}$, under the constraint that for every bidder b_j , $\pi_{b_j^i} \leq u_j^i(\pi_{c_i})$.*

PROOF. The resulting solution is certainly valid: first of all, at least as much money is collected as is given away, because $\sum_{b_j} \pi_{b_j} - \sum_{c_i} \pi_{c_i} = \sum_{b_j} \sum_{c_i} \pi_{b_j^i} - \sum_{c_i} \pi_{c_i} = \sum_{c_i} ((\sum_{b_j} \pi_{b_j^i}) - \pi_{c_i})$ —and the terms of this summation are the objectives of the individual optimization problems, each of which can be set at least to 0 (by setting all the variables are set to 0), so it follows that the expression is nonnegative. Second, no bidder b_j pays more than she is willing to, because $u_j - \pi_{b_j} = \sum_{c_i} u_j^i(\pi_{c_i}) - \sum_{c_i} \pi_{b_j^i} = \sum_{c_i} (u_j^i(\pi_{c_i}) - \pi_{b_j^i})$ —and the terms of this summation are nonnegative by the constraints we imposed on the individual optimization problems.

All that remains to show is that the solution is optimal. Because in an optimal solution, we will extract as much payment from the bidders as possible given the π_{c_i} , all we need to show is that the π_{c_i} are set optimally by this approach. Let $\pi_{c_i}^*$ be the amount paid to charity π_{c_i} in some optimal solution. If we change this amount to π_{c_i}' and leave everything else unchanged, this will only affect the payment that we can extract from the bidders because of this particular charity, and the difference in surplus will be $\sum_{b_j} u_j^i(\pi_{c_i}') - u_j^i(\pi_{c_i}^*) - \pi_{c_i}' + \pi_{c_i}^*$. This expression is, of course, 0 if $\pi_{c_i}' = \pi_{c_i}^*$. But now notice that this expression is maximized as a function of π_{c_i}' by the decomposed solution for this charity (the terms without π_{c_i}' in them do not matter, and of course in the decomposed solution we always set $\pi_{b_j^i} = u_j^i(\pi_{c_i})$). It follows that if we change π_{c_i} to the decomposed solution, the change in surplus will be at least 0 (and the solution will still be valid). Thus, we can change the π_{c_i} one by one to the decomposed solution without ever losing any surplus. \square

THEOREM 4. *When all bids are quasilinear and surplus is the objective, DONATION-CLEARING can be done in linear time.*

PROOF. By Lemma 1, we can solve the problem separately for each charity. For charity c_i , this amounts to maximizing $(\sum_{b_j} u_j^i(\pi_{c_i})) - \pi_{c_i}$ as a function of π_{c_i} . Because all its terms are piecewise linear functions, this whole function is piecewise linear, and must be maximized at one of the points where it is nondifferentiable. It follows that we need only check all the points at which one of the terms is nondifferentiable. \square

Unfortunately, the decomposing lemma does not hold for payment maximization.

PROPOSITION 1. *When the objective is payment maximization, even when bids are quasilinear, the solution obtained by decomposing the problem across charities is in general not optimal (even with concave bids).*

PROOF. Consider a single bidder b_1 placing the following quasilinear bid over two charities c_1 and c_2 : $u_1^1(\pi_{c_1}) = 2\pi_{c_1}$ for $0 \leq \pi_{c_1} \leq 1$, $u_1^1(\pi_{c_1}) = 2 + \frac{\pi_{c_1} - 1}{4}$ otherwise; $u_1^2(\pi_{c_2}) = \frac{\pi_{c_2}}{2}$. The decomposed solution is $\pi_{c_1} = \frac{7}{3}$, $\pi_{c_2} = 0$, for a total donation of $\frac{7}{3}$. But the solution $\pi_{c_1} = 1$, $\pi_{c_2} = 2$ is also valid, for a total donation of $3 > \frac{7}{3}$. \square

In fact, when payment maximization is the objective, DONATION-CLEARING remains (weakly) NP-complete in general. (In the remainder of the paper, proofs are omitted because of space constraint.)

THEOREM 5. *DONATION-CLEARING is (weakly) NP-complete when payment maximization is the objective, even when every bid is concave and quasilinear.*

However, when the bids are also concave, a simple greedy clearing algorithm is optimal.

THEOREM 6. *Given a DONATION-CLEARING instance with payment maximization as the objective where all bids are quasilinear and concave, consider the following algorithm. Start with $\pi_{c_i} = 0$ for all charities. Then, letting*

$$\gamma_{c_i} = \frac{d \sum_{b_j} u_j^i(\pi_{c_i})}{d \pi_{c_i}} \quad (\text{at nondifferentiable points, these derivatives should be taken from the right}),$$

increase $\pi_{c_i^*}$ (where $c_i^* \in \arg \max_{c_i} \gamma_{c_i}$), until either $\gamma_{c_i^*}$ is no longer the highest (in which case, recompute c_i^* and start increasing the corresponding payment), or $\sum_{b_j} u_j = \sum_{c_i} \pi_{c_i}$ and $\gamma_{c_i^*} < 1$. Finally,

$$\text{let } \pi_{b_j} = u_j.$$

(A similar greedy algorithm works when the objective is surplus and the bids are quasilinear and concave, with as only difference that we stop increasing the payments as soon as $\gamma_{c_i^*} < 1$.)

10. INCENTIVE COMPATIBILITY

Up to this point, we have not discussed the bidders' incentives for bidding any particular way. Specifically, the bids may not truthfully reflect the bidders' preferences over charities because a bidder may bid *strategically*, misrepresenting her preferences in order to obtain a result that is better to herself. This means the mechanism is not *strategy-proof*. (We will show some concrete examples of this 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. Such mechanisms (for example, first-price auctions, where winners pay the value of their bids) are typically not strategy-proof: if a bidder reports her true valuation for an outcome, then if this outcome occurs, the payment the bidder will have to make will offset her gains from the outcome completely. Of course, we could try to change the rules of the game—which outcome (payment vector to charities) do we select for which bid vector, and which bidder pays how much—in order to make bidding truthfully beneficial, and to make the outcome better with regard to the bidders' true preferences. This is the field of *mechanism design*. In this section, we will briefly discuss the options that mechanism design provides for the expressive charity donation problem.

10.1 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 in the paper 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 known in economics as the *free rider* problem [14].)

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 $\pi_{c_1} \geq 1$, $u_2^1(\pi_{c_1}) = 0$ 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 represented⁵ 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 $\pi_{c_2} \geq 1$, $u_3^2(\pi_{c_2}) = 0$ 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.

10.2 Mechanism design in the quasilinear setting

There are 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 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 mechanisms [24, 4, 10], or the DAGVA mechanism [6, 1]) which we could apply directly to the expressive charity donation problem. Third, as we saw in Section 9, the clearing problem is much easier in

⁵Formally, 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.)

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.

Before presenting the decomposition result, we introduce some terms from game theory. A *type* is a preference profile that a bidder can have and can report (thus, a type report is a bid). *Incentive compatibility (IC)* means that bidders are best off reporting their preferences truthfully; either regardless of the others’ types (*in dominant strategies*), or in expectation over them (*in Bayes-Nash equilibrium*). *Individual rationality (IR)* means agents are at least as well off participating in the mechanism as not participating; either regardless of the others’ types (*ex-post*), or in expectation over them (*ex-interim*). A mechanism is *budget balanced* if there is no flow of money into or out of the system—in general (*ex-post*), or in expectation over the type reports (*ex-ante*). A mechanism is *efficient* if it (always) produces the efficient allocation of wealth to charities.

THEOREM 7. *Suppose all agents’ preferences are quasilinear. 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 (implementation in dominant strategies or Bayes-Nash equilibrium) and a given notion of individual rationality R (ex post, ex interim, or none), satisfies a certain notion of budget balance (ex post, ex ante, or none), and is ex-post efficient. Then there exists such a mechanism for any number of charities.*

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 if we can do better in the setting of donations to charities.

10.3 Impossibility of efficiency

In this subsection, we show that even in a very restricted setting, and with minimal requirements on IC and IR constraints, it is impossible to create a mechanism that is efficient.

THEOREM 8. *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 piecewise linear functions)).*

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 [14]. 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 [15].

Theorem 7 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 8, 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 2. *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 8, apart from the fact that there are now two charities.)*

11. CONCLUSION

We introduced a bidding language for expressing very general types of matching offers over multiple charities. We formulated the corresponding clearing problem (deciding how much each bidder pays, and how much each charity receives), and showed that it is NP-complete to approximate to any ratio even in very restricted settings. We gave a mixed-integer program formulation of the clearing problem, and showed that for concave bids (where utility functions and payment willingness function are concave), the program reduces to a linear program and can hence be solved in polynomial time. We then showed that the clearing problem for a subclass of concave bids is at least as hard as the decision variant of linear programming, suggesting that we cannot do much better than a linear programming implementation for such bids. Subsequently, we showed that the clearing problem is much easier when bids are quasilinear (where payment willingness functions are linear)—for surplus, the problem decomposes across charities, and for payment maximization, a greedy approach is optimal if the bids are concave (although this latter problem is weakly NP-complete when the bids are not concave). For the quasilinear setting, we studied the mechanism design question of making the bidders report their preferences truthfully rather than strategically. We showed that an ex-post efficient mechanism is impossible even with only one charity and a very restricted class of bids. We also showed that even though the clearing problem decomposes over charities in the quasilinear setting, there may be benefits to linking the charities from a mechanism design standpoint.

There are many directions for future research. One is to build a web-based implementation of the (first-price) mechanism proposed in this paper. Another is to study the computational scalability of our MIP/LP approach. It is also

important to identify other classes of bids (besides concave ones) for which the clearing problem is tractable. Much crucial work remains to be done on the mechanism design problem. Finally, are there good iterative mechanisms for charity donation?⁶

12. REFERENCES

- [1] K. Arrow. The property rights doctrine and demand revelation under incomplete information. In M. Boskin, editor, *Economics and human welfare*. New York Academic Press, 1979.
- [2] L. M. Ausubel and P. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1, 2002. No. 1, Article 1.
- [3] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi-unit combinatorial auctions. In *Theoretical Aspects of Rationality and Knowledge (TARK IX)*, Bloomington, Indiana, USA, 2003.
- [4] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [5] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI-02)*, pages 103–110, Edmonton, Canada, 2002.
- [6] C. d’Aspremont and L. A. Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.
- [7] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1:237–267, 1976.
- [8] D. Goldburg and S. McElligott. Red cross statement on official donation locations. 2001. Press release, http://www.redcross.org/press/disaster/ds_pr/011017legitdonors.html.
- [9] R. Gonen and D. Lehmann. Optimal solutions for multi-unit combinatorial auctions: Branch and bound heuristics. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 13–20, Minneapolis, MN, Oct. 2000.
- [10] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [11] L. Khachiyan. A polynomial algorithm in linear programming. *Soviet Math. Doklady*, 20:191–194, 1979.
- [12] R. Lavi, A. Mu’Alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, 2003.
- [13] D. Lehmann, L. I. O’Callaghan, and Y. Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002. Early version appeared in ACM-EC-99.
- [14] A. Mas-Colell, M. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [15] R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.
- [16] G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, 1999. Section 4, page 11.
- [17] N. Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 1–12, Minneapolis, MN, 2000.
- [18] N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 242–252, Minneapolis, MN, 2000.
- [19] D. C. Parkes. iBundle: An efficient ascending price bundle auction. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 148–157, Denver, CO, Nov. 1999.
- [20] M. H. Rothkopf, A. Pekeć, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.
- [21] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, Jan. 2002. Conference version appeared at the International Joint Conference on Artificial Intelligence (IJCAI), pp. 542–547, Stockholm, Sweden, 1999.
- [22] T. Sandholm, S. Suri, A. Gilpin, and D. Levine. CABOB: A fast optimal algorithm for combinatorial auctions. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1102–1108, Seattle, WA, 2001.
- [23] J. Tagliabue. Global AIDS Funds Is Given Attention, but Not Money. *The New York Times*, June 1, 2003. Reprinted on http://www.healthgap.org/press_releases/a03/060103_NYT_HGAP_G8_fund.html.
- [24] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [25] P. R. Wurman and M. P. Wellman. AkBA: A progressive, anonymous-price combinatorial auction. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 21–29, Minneapolis, MN, Oct. 2000.
- [26] M. 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)*, Acapulco, Mexico, Aug. 2003.

⁶Compare, for example, iterative mechanisms in the combinatorial auction setting [19, 25, 2].