Expressive Markets for Donating to Charities*

Vincent Conitzer
Department of Computer Science
Duke University
Durham, NC 27708, USA
conitzer@cs.duke.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
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 a bargaining chip 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 donor 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 all parties (donors as well as charities) would have preferred 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 cannot be approximated to any ratio in polynomial time unless P=NP, 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. We also consider the case where each charity has a target amount, and bidders' willingness-to-pay functions are concave. Here, we show that the optimal surplus can be approximated to a ratio m, the number of charities, in polynomial

*A short, early conference version (EC-04) of this paper was based on work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IIS-9800994, ITR IIS-0081246, and ITR IIS-0121678. Several sections were added in 2009/2010; this work was supported by NSF grants IIS-0812113, IIS-0905390, and CAREER IIS-0953756, as well as a Research Fellowship from the Alfred P. Sloan Foundation. We are grateful for receiving many useful comments on this work from, among others, the EC and AIJ reviewers as well as participants at the 2004 EC conference, the 2004 World Congress of the Game Theory Society, and the 2005 Dagstuhl seminar on Computing and Markets. These comments, as well as the participation of the Dagstuhl seminar participants in an experiment described in Section 10, have improved the paper significantly.

time (and no significantly better approximation is possible in polynomial time unless P=NP); no polynomial-time approximation ratio is possible for maximizing the total donated, unless P=NP. Subsequently, we show that the clearing problem is much easier when bids are quasilinear—for maximizing surplus, the problem decomposes across charities, and for maximizing the total donated, a greedy approach is optimal if the bids are concave (although this latter problem is weakly NP-hard 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 can be benefits to linking the charities from a mechanism design standpoint. Finally, we discuss an experiment in which we used this methodology to collect money for victims of the 2004 Indian Ocean Tsunami.

1 Introduction

When money is donated to a charitable (or other) cause (hereafter referred to as a *charity*), often the donating party gives *unconditionally*: a fixed amount is transferred from the donor 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 a bargaining chip 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 that 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.

One method that is often used to bring this about 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 then 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 [35].

However, there are several 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 various charities. Moreover, the amount the party offers to donate at different levels of appreciation is allowed to vary arbitrarily. 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. This can be useful in the context of multiple individuals who wish to make matching offers, but the parties do not need to be individuals; for example, one can imagine applying this approach at an international aid conference (for instance, for rebuilding Haiti after its devastating 2010 earthquake) where the parties are donor nations deciding how much each of them will contribute.

However, as we will show, with multiple charities, finding this solution (the *clearing problem*) requires solving an optimization problem that, in general, is hard. A large part of this paper is devoted to studying how hard this problem is under different assumptions on the structure of the offers (or *bids*), and providing algorithms for solving it. Towards the end of the paper, we also study the *mechanism design* problem of motivating the bidders to bid truthfully. We also discuss a small experiment.

In short, expressive markets for making charitable donations have the potential to increase welfare by facilitating the voluntary reallocation of wealth. To reach this potential, we discuss primarily computational aspects, and secondarily mechanism-design aspects of this problem. In Appendix A, we discuss the relationship between expressive charity donation and combinatorial auctions and exchanges. It can safely be skipped, but may be of interest to the reader with a background in combinatorial auctions and exchanges.

[1] Typically, 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 [17].

2 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[2] to contribute. (The bidder's own contribution is also counted in the vector of total 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.) We note that each bidder specifies only *one* total amount that she is willing to give (for each vector of total payments made to the charities), that is, she does not explicitly specify the charities to which her donation will go.[3] 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 $c_1, c_2, \ldots, c_m$, a* bid *by bidder b_j is a function $w_j : \mathbb{R}^m \to \mathbb{R}$. 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) $w_j(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m})$.*

We note that w_j does *not* necessarily decompose across charities, that is, we do *not* necessarily have that $w_j(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m}) = \sum_{i=1}^{m} w_j^i(\pi_{c_i})$ (for some component functions w_j^i). There are several reasons why such a decomposability assumption may be too restrictive. One is that the charities may be related: for example, two charities may pursue the same goal or related goals. Another reason is that a bidder may only have a limited amount of money to give—or, more generally, a bidder who has already given a large amount may become more reluctant to give another dollar. Later in this paper, we will effectively mostly assume away the former reason; we will discuss the latter reason in more detail shortly.

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 $(\pi_{b_1}, \pi_{b_2}, \ldots, \pi_{b_n})$, and a vector of payments received by the charities $(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m})$.*

Definition 3 *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 $1 \leq j \leq n$, $\pi_{b_j} \leq w_j(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m})$ (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.

[2]The word "willing" here should not be interpreted as being necessarily directly related to the bidder's true preferences. Rather, the bidder just indicates how much she *agrees* to contribute in each case, and at this point we are not yet concerned with the bidder's intentions behind entering into this agreement.

[3]In Appendix B, we discuss a variant where bidders make payments to charities directly and can express that they are not willing to give to certain charities.

Definition 4 *An* objective *is a function from the set of all outcomes to $\mathbb{R}$.*[4] *After all the bids have been collected, a valid outcome will be chosen that maximizes this objective.*

One example of an objective is *surplus*, given by $\sum_{j=1}^{n} \pi_{b_j} - \sum_{i=1}^{m} \pi_{c_i}$. One easy way of thinking about the surplus is as the profit of a company managing the expressive donation marketplace. However, it should be emphasized most strongly that the surplus does not *need* to go the entity organizing the donation market. The surplus can also be returned to the bidders, or to the charities. Indeed, the idea of a company profiting from charitable donations may be unpalatable, and the organization would presumably be run more naturally as a nonprofit.

To illustrate this point more precisely, we can consider the following objective, which we will call *surplus*$'$: $\sum_{j=1}^{n} [w_j(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m}) - \pi_{b_j}]$. That is, for each bidder there is slack between how much she is willing to pay, and how much she actually pays. A large slack indicates a happy bidder: she would have been willing to pay even more to achieve these donations to the charities. (We will give a more precise utility-theoretic interpretation of this shortly.) So, we try to maximize the sum of these slacks; this is the surplus$'$ objective. (We note that this objective depends on the bids.) The next proposition shows that surplus and surplus$'$ are effectively equivalent. As a result, in most of the paper we will simply consider maximizing surplus, with the understanding that such techniques apply just as well to maximizing surplus$'$.

Proposition 1 *A vector $\pi_{c_1}, \ldots, \pi_{c_m}$ is part of a valid outcome maximizing surplus if and only if it is part of a valid outcome maximizing surplus$'$.*

Proof: Suppose the outcome $\pi_{b_1}, \ldots, \pi_{b_n}, \pi_{c_1}, \ldots, \pi_{c_m}$ maximizes surplus among valid outcomes; let $s = \sum_{j=1}^{n} \pi_{b_j} - \sum_{i=1}^{m} \pi_{c_i}$ be the surplus. Then, the outcome where we give the entire surplus to bidder 1, $\pi'_{b_1} = \pi_{b_1} - s, \pi'_{b_2} = \pi_{b_2}, \ldots, \pi'_{b_n} = \pi_{b_n}, \pi_{c_1}, \ldots, \pi_{c_m}$, is still valid, and its surplus$'$ is at least s. Thus, the optimal surplus$'$ is always at least the optimal surplus.[5]

Conversely, suppose the outcome $\pi'_{b_1}, \ldots, \pi'_{b_n}, \pi_{c_1}, \ldots, \pi_{c_m}$ maximizes surplus$'$ among valid outcomes; let $s' = \sum_{j=1}^{n} [w_j(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m}) - \pi'_{b_j}]$ be the surplus$'$. Then, the outcome where we reduce each bidder's slack to zero, $\pi_{b_1} = w_1(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m}), \ldots, \pi_{b_n} = w_n(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m}), \pi_{c_1}, \ldots, \pi_{c_m}$, is still valid, and its surplus is at least s'. Thus, the optimal surplus is always at least the optimal surplus$'$. It follows that the optimal surplus and the optimal surplus$'$ are always equal. Moreover, they are obtained at the

[4] In general, the objective function may also depend on the bids, but the main objective functions under consideration in this paper—surplus or the total amount donated—do not depend on the bids. The techniques presented in this paper will typically generalize to objectives that take the bids into account directly.

[5] It is good to emphasize that at this point we are not considering fairness (though of course we could spread the surplus more equally) or strategic issues. In fact, in mechanism design, there is a growing literature on the highly nontrivial problem of redistributing surplus payments to the bidders without creating perverse incentives (*e.g.*, [4, 27, 5, 19, 23, 1, 6, 20]).

same vectors $\pi_{c_1}, \ldots, \pi_{c_m}$, as shown above. $\square$

While the surplus objective has appealing properties, we may consider others. For example, another objective is *total amount donated*, given by $\sum_{i=1}^{m} \pi_{c_i}$. This objective has the advantage that more money is donated to the charities overall; it also has the downside that more may be given to a charity than it really needs. This distinction between maximizing surplus and maximizing the total amount donated has an analogue in the combinatorial exchanges literature, where two common objectives are to maximize surplus, and to maximize total trading volume (aka. liquidity)—though at least in that context there are good arguments for preferring the surplus objective [34, 16].

Finding a 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 5 (DONATION-CLEARING) *We are given a set of n bids over charities $c_1, c_2, \ldots, c_m$. Additionally, we are given an objective function. We are asked to find an objective-maximizing valid outcome.*

How hard the DONATION-CLEARING problem is depends on the types of bids used and the language in which they are expressed.

2.1 The special case of a single charity

To build intuition, it is helpful to consider an important special case that is interesting in its own right: the case in which there is only one charity, c. In this context, each bidder b_j specifies a function $w_j : \mathbb{R} \to \mathbb{R}$, indicating how much the bidder is willing to give as a function of the total received by the charity. In this special case, the clearing problem is particularly easy. We consider the function $w_{\text{total}} : \mathbb{R} \to \mathbb{R}$, defined by $w_{\text{total}}(\pi_c) = \sum_{j=1}^{n} w_j(\pi_c)$. This function returns, as a function of the amount received by the charity, the total that the bidders are willing to give. Given this function, it is easy to see which amounts the charity can receive in valid outcomes: any amount π_c such that $w_{\text{total}}(\pi_c) \geq \pi_c$ is part of a valid outcome, because if we collect the full amount that each bidder is willing to pay, we will collect at least as much money as is given away. (Of course, we can also collect less, as long as the total collected is at least π_c.) Thus, if we graph the w_{total} function with the total received, π_c, on the x-axis, and the total willingness to pay, $w_{\text{total}}(\pi_c)$, on the y-axis, then we find valid outcomes wherever the function w_{total} is at or above the 45-degree line (the identity function); the outcome that maximizes the total amount donated is the furthest-to-the-right such point, and the outcome that maximizes surplus is the point that is furthest above the 45-degree line. This is illustrated in more detail in the Indian Ocean Tsunami experiment in Section 10.

When there are multiple charities, the clearing problem becomes more complicated. In fact, for multiple charities, it becomes less reasonable to ask bidders to specify arbitrary functions $w_j : \mathbb{R}^m \to \mathbb{R}$, and so we need to think about designing a bidding language for the bidders. Before we do so, it is helpful to discuss whether it makes sense to interpret the bids as statements about the bidders' utilities, and how to do so.

2.2 Discussion of the merits of utility-theoretic interpretations

So far, we have not said anything about how we assume bids relate to the bidders' utilities. At some level, it is not strictly necessary to do so: the semantics of a bid in our framework are perfectly well defined even without any utility-theoretic interpretation. Namely, the payment willingness function simply specifies the maximum amount that a bidder can be asked to pay, given a vector of donations. The bid represents a conditional commitment to donate money. Indeed, in order to apply a system like the one described in this paper in practice, it is highly desirable to have such a simple description of the semantics of bids—one that does not depend on abstract concepts such as utility, which are presumably foreign to many of the people who might use the system.

Moreover, as we will see, even if we do interpret the bids in a utility-theoretic way, the basic design that we propose (when interpreted straightforwardly, that is, as a first-price mechanism) is in any case not incentive compatible. (In mechanism design, a mechanism is *incentive compatible* if each bidder is always best off declaring her preferences truthfully.) To make a comparison to auctions, our proposed design is more similar to a first-price auction, in which the winning bidder simply pays her bid, than to a Vickrey auction [36], in which the winner pays the bid of the next-highest bidder (or to the Vickrey auction's generalization to VCG mechanisms [36, 7, 18]). VCG mechanisms are in fact incentive compatible, and because of this, the bids can arguably truly be interpreted as reflecting the bidders' utilities. In contrast, in first-price auctions (or many other mechanisms), bidders are incentivized to bid strategically. Hence, while their bids may be given in the same form as under a VCG mechanism, it is, to say the least, a stretch to interpret these bids as truly representing the bidders' utilities. Because of this, it is perhaps best to consider a first-price auction as just a particular game—one in which we do not attach too much direct meaning to the bids, other than the guarantee to each bidder that she will not be made to pay more than she bid. It is the same for our charity market.

In spite of the theoretical advantages of incentive compatibility, it is extremely rare to see an incentive compatible mechanism such as VCG actually deployed in practice. This is in contrast to first-price mechanisms, which are quite common. Presumably, one important reason for this is that it is much easier to explain a first-price mechanism to a novice. Other practical drawbacks of incentive compatible mechanisms such as VCG have received much discussion [29, 30, 11, 3]. Moreover, as we will see in Section 9, in the specific context of charitable donations, there are fundamental limitations on what can be achieved by incentive compatible mechanisms.

One may argue that, with a mechanism that is not incentive compatible, we need to be concerned about strategic behavior by the bidders. In fact, by the *revelation principle* from the theory of mechanism design, focusing on incentive compatible mechanisms is without loss of optimality in the context of strategic bidders that behave according to the laws of game theory. This is absolutely a valid point, and indeed we do consider such strategic behavior in Section 9. Nevertheless, *in the specific context of donations to charities*, it seems perhaps unlikely that a game-theoretic solution based on a simple model of utilities will give us an accurate prediction of actual behavior in practice. This is first because it appears that people's reasons for giving to charity are complex and varied. Second, the image of a hard-nosed agent strategically pursuing maximum advantage for herself seems somewhat out of place in the context of charitable donations. This is not a normative statement—in principle, an agent's utility function can model all sorts of preferences, including altruistic ones, and there is no reason that she should not act in accordance with her preferences—but these types of considerations are nevertheless likely to affect people's behavior in practice. To illustrate these points, it is helpful to point at the real-world bids in the experiment in Section 10. It seems difficult to explain such a diversity of bids with a simple strategic model.

In spite of all these arguments, we *do* believe that it is important to think about how bids may reflect the bidders' underlying utility functions. Besides considerations of strategic bidding, one important reason for this is that we need a bidding language in which the bidders can express their bids. A good bidding language makes it easy to express "natural" bids, and in order to understand what natural bids are, it helps to think about how they may relate to the bidder's utility.

2.3 A utility-theoretic interpretation

We now give a utility-theoretic interpretation of bids in our framework. We assume that bidder b_j's utility u_j depends only on how much she gives (π_{b_j}) and how much the charities receive ($\pi_{c_1}, \ldots, \pi_{c_m}$). We additionally assume that the utility is nonincreasing in π_{b_j} (when holding the π_{c_i} fixed), and moreover that for any values of $\pi_{c_1}, \ldots, \pi_{c_m}$, there exists a value of π_{b_j} such that $u_j(\pi_{b_j}, \pi_{c_1}, \ldots, \pi_{c_m}) = 0$. Given this, we can choose to interpret the bidder's willingness to pay as the largest amount that the bidder could pay and still end up with nonnegative utility, that is, $w_j(\pi_{c_1}, \ldots, \pi_{c_m}) = \max\{\pi_{b_j} | u_j(\pi_{b_j}, \pi_{c_1}, \ldots, \pi_{c_m}) \geq 0\}$.

Given this, the condition that a bidder should not pay more than $w_j(\pi_{c_1}, \ldots, \pi_{c_m})$ is equivalent to saying that she should receive utility at least 0. This can be interpreted as an *individual rationality* (aka. *voluntary participation*) constraint: participating in the market should not make a bidder worse off than she would have been if the whole event had not happened and no bidder had given anything.[6]

An interesting special case is that of *quasilinear utility*, where we can write $u_j(\pi_{b_j}, \pi_{c_1}, \ldots, \pi_{c_m}) = w_j(\pi_{c_1}, \ldots, \pi_{c_m}) - \pi_{b_j}$. In this case, utility can be expressed in monetary terms. This also gives us another justification for maximizing surplus (more precisely, surplus$'$):

Proposition 2 *If utilities are quasilinear, then an outcome maximizes surplus$'$ if and only if it maximizes the sum of the bidders' utilities.*

[6]One might argue that a more appropriate definition of individual rationality would be that a bidder should not be better off acting separately in the world, while the other bidders continue to participate in the mechanism. However, such a more stringent definition immediately results in impossibilities. For example, consider a charity that requires a donation of 1 for a project (but has no use for additional money). If there is only a single bidder, who is willing to pay 2 to see the charity get 1, then presumably this bidder should indeed give 1—in fact, the more stringent definition of individual rationality would require this, because otherwise the bidder could just give the money to the charity directly, outside of the mechanism. But then, consider the situation where there are two such bidders. If one of them does not participate, then the other bidder will end up giving the full amount, so that the former bidder free-rides. But we cannot make *both* bidders at least as happy as they would be as free-riders—someone has to put up the money. Thus, there is no individually rational outcome at all under this more stringent definition. Section 9 considers strategic phenomena in more detail.

Proof: A bidder's utility for an outcome is $w_j(\pi_{c_1}, \ldots, \pi_{c_m}) - \pi_{b_j}$, which is equal to that bidder's slack—and the surplus$'$ objective is the sum of these slacks. □

While we will devote a reasonably large part of this paper to the case of quasilinear utilities, in our opinion, this restriction is often not reasonable in the context of donations to charities. Intuitively, the argument is as follows. If utility were truly quasilinear, then the reason that a donor gives only a bounded amount to a charity must be that the donor feels that, eventually, the marginal benefit of more money going to the charity starts to significantly decrease. While this may be reasonable for small charities with a limited mission (*e.g.*, the local animal shelter only needs so much money to do the most important things that it can), for large charities with a broader mission (*e.g.*, fighting world hunger), it seems that the marginal benefit of more money going to the charity stays very nearly constant over a very large range of amounts of money donated. In our opinion, in this context it is more reasonable to argue that the marginal utility that the donor has for keeping money to herself changes: she can easily spare a small amount of money, but to give a medium-sized amount of money, she has to give up some luxury goods that she likes, and to give a large amount, she has to give up more essential goods. This is the reason that she stops giving after some point. That is, her utility for money is strictly concave, not linear.[7] (It should be emphasized that we only really wish to argue that this is *one* possible reason for donors giving a bounded amount, and therefore that our model should be able to accommodate it. In fact, our model also accommodates other reasons for giving a bounded amount, such as the charity only needing a certain amount to do the most important things.)

3 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 $\{0, 1, 2, \ldots, s\}$, this would still require a bidder to specify $(s + 1)^m$ values, which is exponential in the number of charities. Thus, in the context of multiple charities, we need a bidding language that will allow the bidders to at least specify *some* bids more concisely. We will introduce a bidding language that only represents a subset of all possible bids, which can be described concisely.[8] On the one hand, we believe that this language allows bidders to represent bids that are natural and useful. On the other hand, we do not intend for this to be the final word on bidding languages in this domain: certainly, other languages can be created that allow the bidders to express different bids that may still be natural and useful. Nevertheless, we consider it likely that the results obtained for our language will either generalize directly to such languages, or at least provide a useful starting point.

[7] This argument does presuppose that the donor has very pure motivations, that is, the donor's reason for giving is purely that she wants to see the charity receive money. This is in contrast to, for example, a donor who feels that there is a social expectation from her friends, or an ethical obligation on her own part, to give a particular amount.

[8] Of 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.

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 *valuation function* $v_j : \mathbb{R}^m \to \mathbb{R}$, specifying how much bidder j "appreciates" a given vector of total donations to the charities. Then, there is a *donation willingness* or *willingness-to-pay* function $w_j : \mathbb{R} \to \mathbb{R}$, which specifies how much bidder j is willing to pay given her valuation $v_j(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m})$ for the vector of donations to the charities. Note that we are overloading notation here: before, the domain of w_j functions consisted of vectors of total amounts donated to the charities, but here, the domain is that of nonnegative real numbers, representing the bidder's total valuation (which is itself a function of the total amounts donated to the charities). We overload notation because the range of the function is the same in both cases: nonnegative real numbers, representing how much the bidder is willing to pay. The primary use of w_j in the rest of this paper is as a mapping from valuations to willingness to pay. We emphasize that this function does *not* need to be linear, so that valuations should not be thought of as necessarily expressible in dollar amounts. (Indeed, we argued above that 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, relating the two different uses of w_j, we have $w_j(v_j(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m})) = w_j(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m})$, and we let the bidder describe her functions v_j and w_j separately; she will submit these functions as her bid.

Our first restriction is that the valuation that a bidder derives from money donated to one charity is *independent* of the amount donated to another charity. Thus, $v_j(\pi_{c_1}, \pi_{c_2}, \ldots, \pi_{c_m}) = \sum_{i=1}^{m} v_j^i(\pi_{c_i})$. (We observe that this does *not* imply that the entire bid decomposes similarly, because of the possible nonlinearity of w_j.) This restriction seems reasonable in many cases, though it may be called into question in cases where multiple charities are working on similar or related projects (for example, disaster relief). If multiple charities are working on the *same* project, then we can simply consider them a single charity (in fact, this is what we did in the tsunami event described in Section 10).

Furthermore, for concreteness, we require that each v_j^i be piecewise linear. An interesting special case that we will study is when each v_j^i is a line: $v_j^i(\pi_{c_i}) = a_j^i \pi_{c_i}$. 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 value of a dollar to the charity is not (significantly) affected by the amount given by the bidders.

The only restriction that we place on the bidders' payment willingness functions w_j is that they be piecewise linear. One interesting special case is a *threshold bid*, where w_j is a step function: the bidder is willing to give t dollars if her valuation exceeds s, and otherwise 0. Another interesting special case is when such a bid is *partially acceptable*: the bidder is willing to give t dollars if her valuation exceeds s; but if her valuation is $v < s$, she is still willing to give $\frac{vt}{s}$ dollars. We are interested in such simple bids on the one hand for technical reasons, for example, showing hardness results even for step functions; on the other hand, they do appear to come up in practice—for example, five of the eighteen bids in the experiment in Section 10 were threshold bids.

One might wonder why, if we are given the bidders' valuation functions, we do not simply maximize the sum of the valuations, rather than surplus or total donated. There are several reasons. First, it would be possible for a bidder to inflate her valuation by changing its units (and correspondingly changing her willingness to pay to adjust for the change in units), thereby making her bid more important for valuation maximization purposes even though the bid has effectively remained the same. Second, a bidder could simply give a payment willingness function that is 0 everywhere, and have her valuation be taken into account in deciding on the outcome, in spite of her not contributing anything. We also recall that Proposition 2 states that maximizing surplus$'$ maximizes the sum of the bidders' utilities if utilities are quasilinear.

We now give an example that illustrates what bids may look like in this framework, as well as what outcomes may result under the different objectives.

Example 1 *Let us consider an example with two bidders and two charities. The first bidder bids:*

- $v_1^1(\pi_{c_1}) = (3/4)\pi_{c_1}$ *for* $\pi_{c_1} < 1$, $v_1^1(\pi_{c_1}) = 3/4$ *for* $\pi_{c_1} \geq 1$ *(valuation function for the first charity);*
- $v_1^2(\pi_{c_2}) = (1/2)\pi_{c_2}$ *for* $\pi_{c_2} < 1$, $v_1^2(\pi_{c_2}) = 1/2$ *for* $\pi_{c_2} \geq 1$ *(valuation function for the second charity);*
- $w_1(v_1) = v_1$ *(payment willingness function).*

The second bidder bids:

- $v_2^1(\pi_{c_1}) = 0$ *for* $\pi_{c_1} < 1$, $v_2^1(\pi_{c_1}) = 1/2$ *for* $\pi_{c_1} \geq 1$ *(valuation function for the first charity);*
- $v_2^2(\pi_{c_2}) = (1/4)\pi_{c_2}$ *for* $\pi_{c_2} < 1$, $v_2^2(\pi_{c_2}) = 1/4 + (1/8)(\pi_{c_2} - 1)$ *for* $\pi_{c_2} \geq 1$ *(valuation function for the second charity);*
- $w_2(v_2) = v_2$ *for* $v_2 < 1/2$, $w_2(v_2) = 1/2 + (1/2)(v_2 - 1/2)$ *for* $v_2 \geq 1/2$ *(payment willingness function).*

The valid outcome that maximizes surplus is $\pi_{b_1} = 3/4, \pi_{b_2} = 1/2, \pi_{c_1} = 1, \pi_{c_2} = 0$*, for a surplus of* $1/4$*. The valid outcome that maximizes the total donated is* $\pi_{b_1} = 13/12, \pi_{b_2} = 7/12, \pi_{c_1} = 1, \pi_{c_2} = 2/3$*. We note that, intuitively, under the surplus objective, we carefully evaluate whether additional donations to the charities are worthwhile,*[9] *whereas under the objective of maximizing the total donated, we continue to give to the charities in a way that maximizes the additional willingness to donate, until the bidders' willingness to pay is entirely exhausted.*

[9]For clarity, we emphasize that this evaluation is done taking the bids at face value, that is, we do not consider here whether the bidders have perhaps misreported their preferences and try to assess whether additional donations are worthwhile with respect to the *true* preferences.

4 Hardness of clearing the market

In this section, we will show that the clearing problem is completely inapproximable, even when every bidder's valuation 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 [14].

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 v_j^i is a line; that is, the valuation that any bidder derives from any charity is linear.*
2. *All the v_j^i have slope either 0 or 1.*
3. *Every bidder either has at most 2 charities that affect her valuation (with slope 1), or all charities affect her valuation (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: In this proof, we will represent bids as follows: $(\{(c_1, a_1), (c_2, a_2), \ldots\}, s, t)$ indicates that $v_j^k(\pi_{c_k}) = a_k \pi_{c_k}$ (this function is 0 for charities c_k not mentioned in the bid), and $w_j(v_j) = t$ for $v_j \geq s$, $w_j(v_j) = 0$ otherwise. We say a bid b is *accepted* if its threshold s_b is reached, in which case we will have $\pi_b = t_b$.

We reduce an arbitrary MAX2SAT instance, given by a set of clauses $K = \{k_1, k_2, \ldots\} = \{\{l_{k_1}^1, l_{k_1}^2\}, \{l_{k_2}^1, l_{k_2}^2\}, \ldots\}$ (where l_k^1, l_k^2 are the literals in clause k) over a set of variables 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 $b_v = (\{(c_{+v}, 1), (c_{-v}, 1)\}, 2, 1 - \frac{1}{4|V|})$. ($+v$ and $-v$ are the positive and negative literals corresponding to 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: $b_0 = (\{(c_1, 1), (c_2, 1), \ldots, (c_m, 1)\}, 2|V| + \frac{T}{8|V||K|}, \frac{1}{4} + \frac{1}{16|V||K|})$. We now show that the two instances are equivalent.

First, suppose there exists a solution to the MAX2SAT instance. If in this solution, l is *true*, then let $\pi_{c_l} = 2 + \frac{T}{8|V|^2|K|}$; 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^i) is *true* (that is, k is satisfied) and $c_{l_k^i}$ 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 collected 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 valid 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 at least another 1 from other bids. 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 collected from these bids, so these bids cannot help close the gap. Finally, all the other bids combined can contribute at most $|K|\frac{1}{8|V||K|} + \frac{1}{4} + \frac{1}{16|V||K|} < 1$. Hence, at most one of b_{+v} and b_{-v} is accepted. It follows that we can interpret a valid 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 must satisfy 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, it follows that $b_{l_k^1}$ is not accepted, and as a result 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 $c_{-v_k^1}$. (This follows from the following reasoning. If $c_{-l_k^1}$, where $-l_k^1$ is the negation of l_k^1, receives at least 2, then the total paid to these charities is at least 3 but these bids contribute at most 2; otherwise, if $c_{+v_k^1}$ and $c_{-v_k^1}$ together receive at least 2, then at most 1 can be collected from these bids; otherwise, nothing can be collected from these bids at all.) But this is the same situation that we analyzed before, and we know it is impossible because the other bids cannot close a gap of 1. All that remains to show is that at least T clause bids are accepted, because if so then it follows that our partial assignment satisfies at least T clauses.

We first 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 of the 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, $b_{v_k^1}, b_{+v_k^1}, b_{-v_k^1}$ contribute (at least) $\frac{1}{4|V|}$ less than the total payments to c_{+v} and c_{-v} (because these total payments must be at least 2, and at most one of $b_{+v_k^1}$ and $b_{-v_k^1}$ can be accepted, so the total collected from these three bids is at most $2 - \frac{1}{4|V|}$). 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 must be 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. □

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 positive ratio. This holds even if the DONATION-CLEARING problem instances 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 (that achieves the target number of satisfied clauses), and a zero solution otherwise. So we could decide whether arbitrary MAX2SAT instances have solutions this way, and it would follow that P=NP. □

This should not be interpreted to mean that our approach to the problem of donating to charities is infeasible. First, as we will show, there are very expressive families of bids for which the problem is solvable in polynomial time. (We have also already discussed the case where there is only one charity.) Second, NP-hardness is often overcome in practice (especially when the stakes are high). For instance, even though the problem of clearing combinatorial auctions is NP-hard [28] (even to approximate [31, 37]), it is typically solved to optimality in practice [32].

5 Mixed integer programming formulation

In this section, we give a mixed integer programming (MIP) formulation for the general problem. We also discuss a special case in which this formulation reduces to a linear programming (LP) formulation. In this case, the problem is solvable in polynomial time, because linear programs can be solved in polynomial time [21].

The variables of the MIP that determine the final outcome are the payments made to the charities, denoted by π_{c_i}, and the payments collected from the bidders, π_{b_j}. The objectives we 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 $\sum_{i=1}^{m} \pi_{c_i}$.

The constraint that the outcome should not result in a 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 valuation variable v_j^i indicating the valuation that this bidder derives from the payment to this charity. The bidder's total valuation is given by another variable v_j, with the constraint that $v_j = \sum_{i=1}^{m} v_j^i$.

Each v_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 v_j^i, so that it cannot exceed the given function. The MIP solver can then push the v_j^i variable all the way up to the constraint, in order to collect as much payment from this bidder as possible. In the case where the v_j^i are concave, this is easy: if $(s_l^{i,j}, t_l^{i,j})$ and $(s_{l+1}^{i,j}, t_{l+1}^{i,j})$ are the endpoints of a finite linear segment in the function, we add the constraint that $v_j^i \leq t_l^{i,j} + \frac{\pi_{c_i} - s_l^{i,j}}{s_{l+1}^{i,j} - s_l^{i,j}}(t_{l+1}^{i,j} - t_l^{i,j})$. If the final (infinite) segment starts at $(s_k^{i,j}, t_k^{i,j})$ and has slope d, we add the constraint that $v_j^i \leq t_k^{i,j} + d(\pi_{c_i} - s_k^{i,j})$. Using the fact that the function is concave, for each value of π_{c_i}, the tightest upper bound on v_j^i is the one corresponding to the segment corresponding to that value of π_{c_i}, and therefore these constraints are sufficient to get the correct value of v_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: $(s_{k+1}^{i,j}, t_{k+1}^{i,j}) = (s_k^{i,j} + M, t_k^{i,j} + dM)$, 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 $x_l^{i,j} \in \{0, 1\}$ be an indicator variable that should be 1 if π_{c_i} corresponds to the lth segment of the function, and 0 otherwise. To ensure 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 $\pi_{c_i} = \sum_{l=0}^{k+1} s_l^{i,j} \lambda_l^{i,j}$, so that the $\lambda_l^{i,j}$ must be set correctly. Then, we can set $v_j^i = \sum_{l=0}^{k+1} t_l^{i,j} \lambda_l^{i,j}$. (This is a standard MIP technique [25].)

Finally, each π_{b_j} is bounded by a function of v_j: the (piecewise linear) function w_j provided by the bidder. Representing this function is entirely analogous to how we represented v_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 valuation function v_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 v_j^i and w_j are concave (and piecewise linear), the DONATION-CLEARING problem can be solved in polynomial time using linear programming.*

6 Why one cannot do much better than linear programming

One may wonder whether, for the special case of the DONATION-CLEARING problem presented in Theorem 2 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 by Theorem 2). (The decision variant of a maximization 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 about as fast. (And thus, we could solve the optimization version of the linear program with binary search.)

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

Theorem 3 *The question of whether an LP (given by a set of linear constraints*[10]*) has a feasible solution can be modeled as a DONATION-CLEARING instance with maximizing the total donated 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 instance, each bid b_j has only linear v_j^i functions, and is a partially acceptable threshold bid ($w_j(v_j) = t_j$ for $v_j \geq s_j$, otherwise $w_j(v_j) = \frac{v_j t_j}{s_j}$). 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 this proof, we will represent bids as follows: $(\{(c_1, a_1), (c_2, a_2), \ldots\}, s, t)$ indicates that $v_j^k(\pi_{c_k}) = a_k \pi_{c_k}$ (this function is 0 for c_k not mentioned in the bid), and $w_j(v_j) = t$ for $v_j \geq s$, $w_j(v_j) = \frac{v_j t}{s}$ 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

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

$b_j = (\{(c_{-x_i}, r_i^j)\}_{i:r_i^j>0} \cup \{(c_{+x_i}, -r_i^j)\}_{i:r_i^j<0}, (\sum_i |r_i^j|)H - s_j, 1)$. Let the target total amount donated be $2vH$.

Suppose there is a feasible solution $(x_1^*, x_2^*, \ldots, x_v^*)$ to the LP. Without loss of generality, we can suppose that $|x_i^*| \leq H$ for all i. Then, in the DONATION-CLEARING instance, 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 collect the maximum payment from the bids b_{x_i}—a total payment of $2vH - c$. Additionally, the valuation 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 that we can collect the maximum payment from all the bids b_j, for a total payment of c. It follows that we can collect 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 $2vH$. Then the maximum payment must be collected from each bidder. From the fact that the maximum payment must be collected 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 total payment that can be collected is $2vH$, it follows that for each i, $\pi_{c_{+x_i}} + \pi_{c_{-x_i}} = 2H$ exactly. Let $x_i^* = \pi_{c_{+x_i}} - H = H - \pi_{c_{-x_i}}$. Then, from the fact that the maximum payment must be collected from each bid 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. $\square$

7 Target amounts and concave payment willingness functions

In this section, we study a special case of the DONATION-CLEARING problem where the following two conditions hold. First, every charity c_i has a target amount τ_{c_i} that it is seeking to collect. If less money than this is collected, it is useless to the charity; if any additional money beyond the target amount is collected, then this additional money is also useless. (For example, the charity may be completely devoted to the project of drilling a water well for a particular community, which will cost a fixed amount.) The target amounts are common knowledge, and as a result, bidder b_j derives a fixed amount of valuation υ_{b_j,c_i} from charity c_i if and only if the charity has achieved its target amount. That is, $v_j^i(\pi_{c_i}) = \upsilon_{b_j,c_i}$ if $\pi_{c_i} \geq \tau_{c_i}$, and $v_j^i(\pi_{c_i}) = 0$ otherwise. Additionally, we assume that the payment willingness functions w_j are concave. Under these conditions, we will show that, when the objective is to maximize surplus, the DONATION-CLEARING problem can be approximated in polynomial time to a ratio of m, the number of charities—still not a very positive result. However, we will also show that no significantly better result is possible unless P=NP. For maximizing the total amount donated, we will show that no positive approximation at all is possible unless P=NP.

Consider the following Greedy algorithm for DONATION-CLEARING in this context. We start with the outcome where no charity receives any money; we will iteratively decide to give some charities their target amount. At any point, if we take a charity that currently receives nothing, giving that charity its target amount will increase some donors' willingness to pay, resulting in a net effect on surplus. We repeatedly find the charity that results in the greatest increase in surplus, and give it its target amount—until every remaining charity results in a net decrease in surplus, at which point we stop.

Theorem 4 *In the context where every charity c_i has a target amount τ_{c_i}, and bidders' payment willingness functions are concave, the* Greedy *algorithm for DONATION-CLEARING results in an m-approximation to the maximum possible surplus. On the other hand, there are instances where* Greedy *obtains a surplus that is arbitrarily close to $1/(m-1)$ of the maximum possible surplus.*

Proof: We will show that the approximation ratio m is already obtained after the Greedy algorithm has selected its first charity to donate money to; because the surplus obtained by this algorithm never decreases as more charities are selected, this proves the result. Consider a surplus-maximizing solution, with surplus OPT. In such an optimal solution, there is some subset C' of the charities that each receive their target amount (and the others receive nothing). Arbitrarily order the subset C'. If we imagine giving these charities their target amount in sequence, then each charity has a marginal effect on surplus (which depends on the order). The largest of these marginal effects on surplus is at least OPT$/m$; say that the corresponding charity is c_{i^*}. Then, at the beginning of the Greedy algorithm, choosing c_{i^*} must also result in a marginal effect on surplus of at least OPT$/m$. This is because the marginal effect on surplus can only get smaller as more charities have already received their target amounts, because bidders' payment willingness functions are concave. It follows that the first charity chosen by the Greedy algorithm has a marginal effect on surplus of at least OPT$/m$.

To show that the ratio can be as bad as $m-1$, consider the following situation. Each of the charities $c_1, \dots, c_{m-1}$ has a target amount of $1-\epsilon$. Charity c_m has a target amount of $m-1-\epsilon'$ where ϵ' is slightly larger than ϵ. There are $n = m-1$ bidders. Bidder b_j obtains a valuation of 1 if c_j receives its target amount, and also a valuation of 1 if c_m receives its target amount. Bidder b_j's payment willingness function is as follows: $w_j(v_j) = v_j$ for $v_j \leq 1$, and $w_j(v_j) = 1$ for $v_j \geq 1$. The Greedy algorithm will choose to give c_m its target amount first, as this results in a surplus of ϵ'. After this, the algorithm ends because no bidder is willing to pay any more. However, a better solution is to give $c_1, \dots, c_{m-1}$ their target amounts, which results in a surplus of $(m-1)\epsilon$. As ϵ' converges downward to ϵ, we get the desired result. □

Of course, the above approximation ratio is a worst-case result, and it seems that this greedy algorithm is likely to fare much better in practice. Nevertheless, in the worst case, we cannot hope for a significantly better result unless P=NP, as the following result makes clear.

Theorem 5 *In the context where every charity c_i has a target amount τ_{c_i}, and bidders' payment willingness functions are concave, it is not possible to approximate the optimal surplus to a ratio $m^{1-\epsilon}$ in polynomial time, unless P=NP. This is true even if each bidder's willingness-to pay function is linear up to a point and flat after that (when the bidder's budget has been exhausted)—that is, the bid is a partially acceptable threshold bid.*

Proof: We will prove this by reduction from INDEPENDENT-SET, in which we are given a graph and are asked to find a maximum-size set of vertices with no edge between any pair of them. It is known that INDEPENDENT-SET cannot be approximated to a ratio $|V|^{1-\epsilon}$ in polynomial time, unless P=NP [37].

For every vertex v in the graph of the independent set instance, construct a charity c_v. Let ν_v be the number of edges that have v as one of their endpoints. Then, let $\tau_{c_v} = \nu_v - \delta$ (for some small δ) be the target amount for the charity c_v. For every edge e, construct a bidder b_e. If $e = (v, w)$, then b_e receives a valuation of 1 if one of c_v and c_w receives its target amount (and 2 if they both do). For each bidder b_j, the payment willingness function is as follows: $w_j(v_j) = v_j$ for $v_j \leq 1$, and $w_j(v_j) = 1$ for $v_j \geq 1$. Hence, a bidder is willing to give 1 if at least one of its two charities receives its target amount, and 0 otherwise.

We first claim that if δ is sufficiently small, then in any feasible solution, the charities that receive their target amount must correspond to an independent set in the graph. This is because if V' is the set of vertices v so that c_v receives its target amount, and there is an edge between two of the vertices in V', then the number of bidders that are willing to pay an amount of 1 is at most $(\sum_{v \in V'} \nu_v) - 1$. However, the amount that these charities require is $(\sum_{v \in V'} \nu_v) - |V'|\delta$, which is larger for sufficiently small δ—contradicting the supposed feasibility of the solution.

On the other hand, if V' is an independent set, then the number of bidders that are willing to pay an amount of 1 is $\sum_{v \in V'} \nu_v$; the amount that these charities require is $(\sum_{v \in V'} \nu_v) - |V'|\delta$, resulting in a surplus of $|V'|\delta$.

Thus, independent sets correspond exactly to feasible solutions, and the surplus obtained in such a solution is proportional to the size of the independent set, proving the result. □

So far in this section, we have not yet discussed the objective of maximizing the total amount donated. For this, we cannot even obtain an m-approximation in the worst case, as the following result shows.

Theorem 6 *In the context where every charity c_i has a target amount τ_{c_i}, and bidders' payment willingness functions are concave, it is not possible to approximate the maximum amount that can be donated to any ratio $f(m, n) > 0$ in polynomial time, unless P=NP. This is true even if each bidder's willingness-to pay function is linear up to a point and flat after that (when the bidder's budget has been exhausted)—that is, the bid is a partially acceptable threshold bid.*

Proof: From the proof of Theorem 5, we know that in the same context, it is NP-hard to decide whether a surplus of at least K can be achieved. We now show how to reduce an arbitrary instance of this decision variant of the surplus maximization problem to an instance of the problem of maximizing the total donated. To do so, we take any such instance, and leave the existing bidders and charities untouched; we add a single charity c' whose target amount is $\tau_{c'} = H$ for some $H > K$. We also add a single new bidder b', who only cares about charity c'; this bidder receives a valuation of $H - K$ if charity c' achieves its target amount, and is willing to give his valuation, that is, his payment willingness function is the identity function (which can be capped at $H - K$ if desired). None of the other (original) bidders care about charity c'.

Now, if the original instance has a solution with surplus at least K, then in the modified instance, we can let the original bidders pay the same amounts, and let the original charities receive the same amounts. Then, we can give the surplus of at least K to charity c', and let the new bidder b' donate the remaining $H - K$, so that the charity achieves its target amount. This leads to a total donated of at least H.

Conversely, if there is a solution for the modified instance in which c' achieves its target amount of H, then, because b' can pay only $H - K$, the remaining K must come from the original bidders. However, the original bidders obtain all their valuation from the original charities, so it must be the case that this K was simply left over. Hence, if we look at the restriction of the solution to the original bidders and the original charities, this corresponds to a solution with a surplus of at least K.

It follows that c' can receive its target amount H if and only if the original instance has a solution with a surplus of K. By making H sufficiently large, we can make the ratio (in terms of total donated) between any solution in which c' does not receive its target amount and any solution where it does arbitrarily small. Hence, an algorithm that gives any positive approximation ratio based only on n and m can be used to detect whether a surplus of at least K is possible in the original instance—but this is NP-hard. □

8 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 valuation: that is, $w_j = v_j$.[11] In many 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 valuation 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 valuation from more money being donated to that charity. Thus, the bidders will never reach a high enough valuation for their budget constraint (hard or soft) to kick in, 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

[11] Because the units of valuation are arbitrary, we may as well let them correspond exactly to units of money—so we do not need a constant multiplier.

design, which we will discuss in Section 9. 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 try to maximize surplus. The key observation here is that when bids are quasilinear and the objective is surplus, the clearing problem *decomposes* across charities. That is, we can simply optimize for every charity separately; to optimize for one of the charities we only need to know each bidder's valuation function corresponding to that charity. The intuition is as follows: without the quasilinearity assumption, a donation to one charity impacts the other charities via the payment willingness function, but this effect disappears when the payment willingness function is linear—at least for the case of surplus maximization. (We will see shortly that this does not hold for the objective of maximizing the total donated, because, intuitively, for that objective we may wish to transfer surplus generated by one charity to another charity to increase total payments.)

Lemma 1 *Suppose all the 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, $\pi_{b_j^i} \leq v_j^i(\pi_{c_i})$.*

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 outer summation are the objectives of the problem instances for the individual charities, each of which can be set at least to 0 (by setting all the variables to 0), so it follows that the expression is nonnegative. Second, no bidder b_j pays more than she is willing to, because $v_j - \pi_{b_j} = \sum_{c_i} v_j^i(\pi_{c_i}) - \sum_{c_i} \pi_{b_j^i} = \sum_{c_i}(v_j^i(\pi_{c_i}) - \pi_{b_j^i})$, and the terms of this summation are nonnegative by the constraints we imposed on the individual optimization instances.

All that remains to show is that the solution is optimal. Because in an optimal solution, we will collect as much payment from the bidders as possible given the π_{c_i}, all that 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 collect from the bidders because of this particular charity, and the difference in surplus will be $\sum_{b_j} v_j^i(\pi'_{c_i}) - v_j^i(\pi^*_{c_i}) - \pi'_{c_i} + \pi^*_{c_i}$. 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 $\pi_{b_j^i} = v_j^i(\pi_{c_i})$). 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 7 *When all the bids are quasilinear and surplus is the objective, DONATION-CLEARING can be solved in polynomial time.*

Proof: By Lemma 1, we can solve the problem separately for each charity. For charity c_i, this amounts to maximizing $(\sum_{b_j} v_j^i(\pi_{c_i})) - \pi_{c_i}$ 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. □

As we have discussed earlier in the paper, the computationally trivial special case of a single charity is interesting not only as a subroutine for the case of quasilinear bids when surplus is the objective, but also in and of itself—some of the most practical applications of all of this may involve just a single charity. For example, we note that the tsunami event that we describe in Section 10 was a single-charity event.

Unfortunately, the decomposition lemma does not hold for the objective of maximizing the total donated.

Proposition 3 *When the objective is maximizing the total donated, 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: $v_1^1(\pi_{c_1}) = 2\pi_{c_1}$ for $0 \leq \pi_{c_1} \leq 1$, $v_1^1(\pi_{c_1}) = 2 + \frac{\pi_{c_1}-1}{4}$ otherwise; $v_1^2(\pi_{c_2}) = \frac{\pi_{c_2}}{2}$. The decomposed solution is $\pi_{c_1} = \frac{7}{3}$, $\pi_{c_2} = 0$, for a total donated of $\frac{7}{3}$. But the solution $\pi_{c_1} = 1$, $\pi_{c_2} = 2$ is also valid, for a total donation of $3 > \frac{7}{3}$. □

As an aside, the proof of Proposition 3 illustrates that the objective of maximizing the total donated can result in unexpected outcomes: there is only a single bidder in the proof of Proposition 3, and presumably if this bidder had been able to donate to the charities outside of our system, she would have chosen to give 1 to the first charity and 0 to the second. (This is exactly what would have happened under the surplus$'$ objective.) This type of situation may in principle be prevented by changing the individual rationality criterion, saying that instead of guaranteeing that the bidder would not have been better off if nobody had donated at all, we guarantee that the bidder would not have been better off acting separately in the world. (However, such a more stringent individual rationality criterion leads to difficulties, as we discussed when first mentioning individual rationality.)

In fact, when the objective is to maximize the total donated, DONATION-CLEARING remains (weakly) NP-hard in general (when we do not assume bids are concave).

Theorem 8 *DONATION-CLEARING is (weakly) NP-hard when the objective is to maximize the total donated, even when every bid concerns only one charity (and has a step-function valuation function for this charity), and is quasilinear.*

Proof: We reduce an arbitrary KNAPSACK instance (given by m pairs of integers $(k_i, v_i)_{1 \leq i \leq m}$, a cost limit K, and a target value V), to the following DONATION-CLEARING instance. Let there be $m + 1$ charities, $c_0, c_1, \ldots, c_m$. Let there be one quasilinear bidder b_0 bidding $v_0^0(\pi_{c_0}) = 0$ for $0 \leq \pi_{c_0} < 1$, $v_0^0(\pi_{c_0}) = K + 1$ otherwise. Additionally, for each j with $1 \leq j \leq m$, let there be a quasilinear bidder b_j bidding $v_j^j(\pi_{c_j}) = 0$ for $0 \leq \pi_{c_j} < k_i$, $v_j^j(\pi_{c_j}) = \epsilon v_j$ otherwise (where $\epsilon \sum_{1 \leq j \leq m} v_j < 1$). Let the target total amount donated be $K + 1 + \epsilon V$. We now show the two instances are equivalent.

First, suppose there exists a solution to the KNAPSACK instance, that is, a function $f : \{1, \ldots, m\} \to \{0, 1\}$ so that $\sum_{i=1}^{m} f(i)k_i \leq K$ and $\sum_{i=1}^{m} f(i)v_i \geq V$. Then, let $\pi_{c_0} = 1 + \epsilon V + K - \sum_{i=1}^{m} f(i)k_i$, and for $i > 0$, $\pi_{c_i} = f(i)k_i$, for a total donated of $K + 1 + \epsilon V$. Because $1 + \epsilon V + K - \sum_{i=1}^{m} f(i)k_i \geq 1$, b_0's valuation is $K + 1$. For $j > 0$, b_j's valuation is $f(j)\epsilon v_j$, for a total valuation of $\sum_{j=1}^{m} f(j)\epsilon v_j \geq \epsilon V$ for these m bidders. It follows that the total valuation is at least the total amount donated, and so this corresponds to a valid outcome. So there exists a solution to the DONATION-CLEARING instance.

Now suppose there exists a solution to the DONATION-CLEARING instance. Let $f : \{1, \ldots, m\} \to \{0, 1\}$ be given by $f(i) = 0$ if $\pi_{c_i} < k_i$, and $f(i) = 1$ otherwise. Because the total donated is at least $K + 1 + \epsilon V$, and the amount that can be collected from the bidders is at most $K + 1 + \sum_{j=1}^{m} f(j)\epsilon v_j$, it follows that $\sum_{j=1}^{m} f(j)v_j \geq V$. Also, because the total amount donated to charities 1 through m can be at most $K + \epsilon \sum_{1 \leq j \leq m} v_j < K + 1$, it follows that $\sum_{j=1}^{m} f(j)k_j < K + 1$. Because the k_j are integers, this means $\sum_{j=1}^{m} f(j)k_j \leq K$. So there exists a solution to the KNAPSACK instance. □

However, when the bids are also concave, a simple greedy clearing algorithm is optimal. This algorithm works as follows:

- Start with $\pi_{c_i} = 0$ for all charities.
- Let $\gamma_{c_i} = \frac{d \sum_{b_j} v_j^i(\pi_{c_i})}{d\pi_{c_i}}$ (at nondifferentiable points, these derivatives should be taken from the right).
- Let $c_i^* \in \arg\max_{c_i} \gamma_{c_i}$.
- Increase $\pi_{c_i^*}$ until either $\gamma_{c_i^*}$ is no longer the highest (in which case, recompute c_i^* and start increasing the corresponding payment), or $\sum_{b_j} v_j = \sum_{c_i} \pi_{c_i}$ and $\gamma_{c_i^*} < 1$.
- Finally, let $\pi_{b_j} = v_j$.

Theorem 9 *Given a DONATION-CLEARING instance with maximizing the total donated as the objective where all bids are quasilinear and concave, the above greedy algorithm returns an optimal solution.*

Proof: The outcome is valid because everyone pays exactly what she is willing to, and because there is no budget deficit: $\sum_{b_j} \pi_{b_j} = \sum_{b_j} v_j = \sum_{c_i} \pi_{c_i}$. To show optimality, let $\pi^*_{c_i}$ be the amount paid to charity c_i in some optimal solution, and let π'_{c_i} be the amount paid to charity c_i in the solution given by the greedy algorithm. We first observe that it is not possible that for every i, $\pi^*_{c_i} \geq \pi'_{c_i}$ with at least one of these inequalities being strict. This is because at the solution found by the greedy algorithm, $\gamma_{c_i^*}$ is less than 1; hence, using concavity, if $\pi^*_{c_i} > \pi'_{c_i}$, then $\int_{\pi'_{c_i}}^{\pi^*_{c_i}} \gamma_{c_i} d\pi_{c_i} < \pi^*_{c_i} - \pi'_{c_i}$. In other words, the additional payment that needs to be made to the charity is less than the additional payment that can be collected from the bidders because of the additional payment to this charity. Because the surplus at the greedy algorithm's solution is 0, it follows that if for every i, $\pi^*_{c_i} \geq \pi'_{c_i}$ with at least one of these inequalities being strict, then the surplus at the optimal solution would be negative, and hence the solution would not be valid. Thus, either for all i, $\pi^*_{c_i} \leq \pi'_{c_i}$ (but in this case the greedy solution has at least as large a total donated as the optimal solution, and we are done); or there exist i, j such that $\pi^*_{c_i} > \pi'_{c_i}$ but $\pi^*_{c_j} < \pi'_{c_j}$. It cannot be the case that $\gamma_{c_i}(\pi'_{c_i}) > \gamma_{c_j}(\pi^*_{c_j})$, for then the greedy algorithm would have increased π_{c_i} beyond π'_{c_i} before increasing π_{c_j} beyond $\pi^*_{c_j}$. So, $\gamma_{c_i}(\pi'_{c_i}) \leq \gamma_{c_j}(\pi^*_{c_j})$. Because $\pi^*_{c_i} > \pi'_{c_i}$, and using concavity, if we decrease $\pi^*_{c_i}$ and simultaneously increase $\pi^*_{c_j}$ by the same amount, we will not decrease the total payment we can collect from the bidders—while keeping the payment to be made to the charities the same. It follows this cannot make the solution worse or invalid. We can keep doing this until there is no longer a pair i, j such that $\pi^*_{c_i} > \pi'_{c_i}$ but $\pi^*_{c_j} < \pi_{c_j}$, and by the previous we know that then, for all i, $\pi^*_{c_i} \leq \pi'_{c_i}$—and hence the greedy solution is optimal. □

(A similar greedy algorithm works when the objective is surplus and the bids are quasilinear and concave; the only difference is that we stop increasing the payments as soon as $\gamma_{c_i^*} \leq 1$. Of course the result in Theorem 7 is stronger in the sense that it does not require concavity.)

9 Strategic bidding and 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 for herself. This would mean the market mechanism is not *strategy-proof*. (We will show some concrete examples of this shortly.) This is not too surprising, because if we use the methodology described in the paper so far straightforwardly, the resulting mechanism is, in a sense, a *first-price* mechanism, where the mechanism will collect as much payment from a bidder as her bid allows.[12] 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

[12] The surplus′ objective is an exception.

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 each bid vector, and how much does each bidder pay—in order to make bidding truthfully beneficial, and to make the outcome better with regard to the bidders' *true* preferences. This is the subject of *mechanism design*. In this section, we will briefly discuss the options that mechanism design provides for the expressive charity donation problem.

9.1 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 can 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 valuation for the charity of 0), to avoid having to make any payment, while still benefiting from the other bidders' contributions. (This is known as the *free rider* problem [22].)

With multiple charities, another kind of manipulation can occur, where a bidder attempts to steer others' payments towards her preferred charity. For example, suppose that there are two charities, and three bidders. The first bidder bids $v_1^1(\pi_{c_1}) = 1$ if $\pi_{c_1} \geq 1$, $v_1^1(\pi_{c_1}) = 0$ otherwise; $v_1^2(\pi_{c_2}) = 1$ if $\pi_{c_2} \geq 1$, $v_1^2(\pi_{c_2}) = 0$ otherwise; and $w_1(v_1) = v_1$ if $v_1 \leq 1$, $w_1(v_1) = 1 + \frac{1}{100}(v_1 - 1)$ otherwise. Hence, if there were no other bidders, the first bidder would be willing to pay 1 to charity 1, or to charity 2, but not to both. The second bidder bids $v_2^1(\pi_{c_1}) = 1$ if $\pi_{c_1} \geq 1$, $v_2^1(\pi_{c_1}) = 0$ otherwise; $v_2^2(\pi_{c_2}) = 0$ (always); $w_2(v_2) = \frac{1}{4}v_2$ if $v_2 \leq 1$, $w_2(v_2) = \frac{1}{4} + \frac{1}{100}(v_2 - 1)$ otherwise. Hence, if there were no bidders other than bidders 1 and 2, then, regardless of whether the objective is surplus or total donated, charity 1 would receive at least 1, and charity 2 would receive less than 1. Now, the third bidder's *true* preferences are accurately represented (under the utility-theoretic interpretation given earlier in the paper) by the bid $v_3^1(\pi_{c_1}) = 1$ if $\pi_{c_1} \geq 1$, $v_3^1(\pi_{c_1}) = 0$ otherwise; $v_3^2(\pi_{c_2}) = 3$ if $\pi_{c_2} \geq 1$, $v_3^2(\pi_{c_2}) = 0$ otherwise; and $w_3(v_3) = \frac{1}{3}v_3$ if $v_3 \leq 1$, $w_3(v_3) = \frac{1}{3} + \frac{1}{100}(v_3 - 1)$ otherwise. Now, it is straightforward to check that, if the third bidder bids truthfully, then regardless of whether the objective is surplus or total donated, charity 1 will still receive at least 1, and charity 2 will still receive less than 1. The same is true if bidder 3 does not place a bid at all (as in the case of free-rider manipulation); hence bidder 3's valuation/utility will be 1 in this case. But now, if bidder 3 reports $v_3^1(\pi_{c_1}) = 0$ everywhere; $v_3^2(\pi_{c_2}) = 3$ if $\pi_{c_2} \geq 1$, $v_3^2(\pi_{c_2}) = 0$ otherwise (this part of the bid is truthful); and $w_3(v_3) = \frac{1}{3}v_3$ if $v_3 \leq 1$, $w_3(v_3) = \frac{1}{3}$ 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 valuation/utility to bidder 3, it follows that this bidder's 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.

9.2 Mechanism design in the quasilinear setting

There are at least four (interrelated) 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 [36, 7, 18], or the dAGVA mechanisms [13, 2]). Third, as we saw in Section 8, the clearing problem is much easier in the quasilinear 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, it makes sense to pursue social welfare (the sum of the bidders' utilities) as the objective, because here 1) units of valuation correspond directly to units of money, so that we do not have any problem of the bidders arbitrarily scaling their valuations; and 2) it is no longer possible to give a payment willingness function of 0 while still affecting the donations through a valuation function. It is also helpful to recall Proposition 2 here, which states that with quasilinear utilities, an outcome maximizes surplus$'$ if and only if it maximizes the sum of the bidders' utilities.

Before presenting the decomposition result, we introduce some concepts from game theory and mechanism design. A *type* represents particular preferences 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' reported types (*in dominant strategies*), or in expectation over them assuming truthful reporting by the other bidders (*in Bayes-Nash equilibrium*). *Individual rationality (IR)* means bidders are at least as well off participating in the mechanism as not participating; either regardless of the others' reported types (*ex post*), or in expectation over them assuming truthful reporting by the other bidders (*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 (*ex ante*). A mechanism is *efficient* if it (always) produces an efficient allocation—that is, an allocation that maximizes the sum of the bidders' utilities.

Proposition 4 *Suppose all bidders' 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.*

Proof: The mechanism is simply the following: for each charity, run the single-charity mechanism on the bidders' preferences for that charity, and let the bidders make the corresponding payments to that charity. (So, each bidder's total payment will be the sum of her payments to the individual charities.) It is straightforward to check that the desired properties of this combined mechanism follow from the fact that the single-charity mechanism satisfies them. □

9.3 Impossibility of efficiency in the mechanism design context

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 10 *There is no mechanism that 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 valuation functions being step functions or both valuation functions being concave piecewise 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 (v is zero everywhere); otherwise, the bidder derives valuation $\frac{5}{4}$ from the charity getting at least 1, and valuation 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, for the sake of contradiction, that 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 1 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 $\pi_2(H,L) + \pi_2(L,L) \leq 0$. So we have reached the desired contradiction.[13] □

[13] As an alternative proof technique (a proof by computer), we let our automated mechanism design software [8, 9] create an optimal mechanism for the (step-function) instance described in the proof, under the required constraints on the mechanism and with social welfare (counting the payments made) as the objective. The resulting 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 1'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 the automated

The case of step functions in this theorem corresponds exactly to the case of a single, fixed-size, nonrival, nonexcludable public good (the "public good" being that the charity receives the required amount)—for which such an impossibility result is already known [22].[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 [24].

Proposition 4 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 10, no such mechanism exists, there may in fact be a benefit to bringing the charities under the same umbrella. The next proposition shows that this can indeed be the case.

Proposition 5 *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. (This holds even when the conditions are the same as in Theorem 10, 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 10, and her preferences for the second charity correspond to the low type in the proof of Theorem 10. 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 10, where we know that it is impossible to get all of ex-post budget balance, ex-post efficiency, and ex-interim individual 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 valuation 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 charity), which is less than reporting truthfully because ex-post IR holds. Second, if one bidder reports

mechanism design software always finds a 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.

[14]Indeed, the framework in this paper can be used for nonrival, nonexcludable public goods more generally: any cause that benefits from money, and from which all agents derive utility, can be thought of as a charity. For example, the project of buying a coffee machine for the department can be thought of as a charity. The example of drilling a water well mentioned earlier perhaps serves to illustrate that there is no sharp distinction between these interpretations of the framework.

a high valuation 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. $\square$

Of course, Proposition 5 merely gives a very particular example where having a single mechanism for multiple charities can help. This still leaves us far from having a general theory of how to design mechanisms for multiple charities, but it does show that we cannot simply decompose the problem across charities (at least not without some further assumptions).

10 An experiment with real bidders and money

We decided to test the basic framework of this paper at the Dagstuhl Workshop on Computing and Markets (2005). During the presentation of this work at that workshop, we announced that we would conduct a charity drive using the bidding language described in this paper. For simplicity, the event was restricted to a single charity, namely the victims of the December 2004 tsunami that devastated coastal areas in Indonesia, Sri Lanka, India, and Thailand. The objective to be maximized in the event was the total amount donated. Participants were given until 6pm on the day after the presentation to submit their bids. Of the workshop's ≥ 46 participants,[15] 18 submitted valid bids before the deadline (including the two authors of this paper). The bids are shown in Figure 1.

The bids were collected in sequence, and every bidder was able to see all the previous bids. Very quickly, a "target" total amount of US $500 emerged as a focal point, so that many bidders made their donations conditional on at least $500 being collected. (In fact, one bidder strategically made his donation conditional on *exactly* $500 being collected: he did not want the effort to reach $500 to fail on his account, but was not interested in donating if the target amount would also be reached without him.) Once this target amount was in fact reached, the remaining bidders set their sights higher, resulting in a final amount of $700.

Four of the bids were unconditional donations (the same amount given at every total amount collected). Five of the bids were simple threshold bids (0 given below a certain amount collected, some constant amount c at or above it). One bid was the sum of an unconditional donation and a simple threshold bid. The remaining eight bids (including the authors' own bids and one bid that was a copy of one of the authors' bids) were more complex.

Sixteen bids were submitted in US dollars, one in Euros, and one in Canadian dollars. (All bids were converted to US dollars.) Figure 2 shows the total amount bidders were willing to give conditional on each total amount donated (this function is the sum of their individual bid functions); the event cleared at the largest feasible point, that is, the largest point at which the curve intersects the identity function (45 degree line), $700.

[15] The number 46 was obtained by counting the number of faces in the group picture.

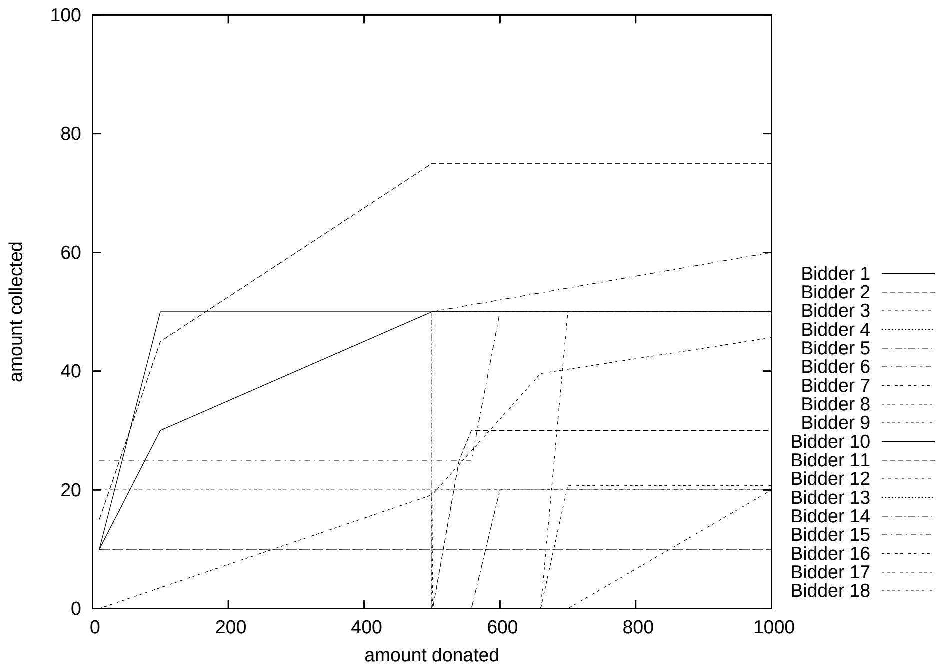

Figure 1: The bids submitted in the tsunami charity event. Each function represents a bidder's willingness to give as a function of the total donated.

From this small experiment, it is clear that at least some bidders prefer to place complex, expressive, conditional bids. Presumably, using this methodology benefited even the bidders who did not make their donation conditional, because their unconditional bids still induced the bidders who did make their bids conditional to give more.

Later, we tried to run a similar event to collect donations for victims of Hurricane Katrina. Unlike the tsunami event, this event was open to everyone and potential participants were approached unsystematically; it was not associated with a workshop or anything of the sort. Unfortunately, the response to this event was rather minimal. While it is not immediately clear which of the various differences between the tsunami and hurricane events were responsible for the difference in success, this perhaps lends some support to the idea that it helps when there are social connections among the bidders in the event. In the next section, we discuss some subsequent research by others that extends our framework with a social-networking component.

11 Subsequent research

Since the conference version of this paper, there have been several highly related works. Most closely related is recent work by Ghosh and Mahdian [15]. In their paper, they take the original version of this paper as a starting point, but then argue that the as-

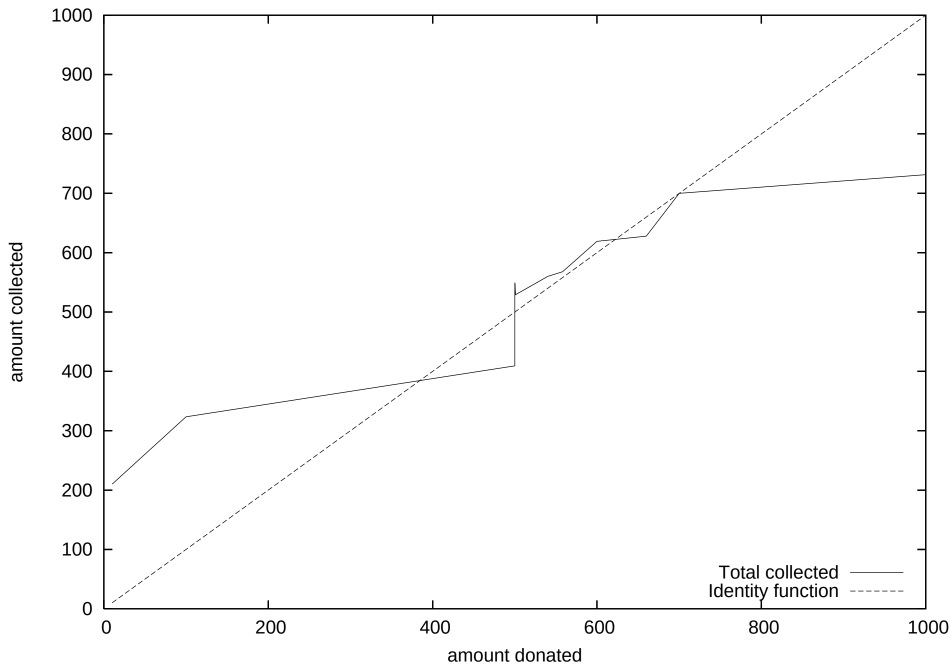

Figure 2: The total willingness to give as a function of the total amount donated (this function is equal to the sum of the bids), and the identity function (the 45 degree line). Since the objective was to maximize the total donated, the event cleared at the largest point at which the curve intersects the identity function, \$700. There is a small spike at \$500 due to the bidder who strategically made his donation conditional on *exactly* \$500 being reached.

sumption that each bidder cares only about the total amount donated to each charity is not realistic. They extend the model so that a bidder can also condition her donation on *who* is making donations. For example, in a social network setting, it may make sense for a bidder to condition her own donation on the donations of her friends only. On the other hand, Ghosh and Mahdian restrict themselves to a single-charity setting, in part because of the hardness results for multiple charities given in our paper. They show that a unique maximal payment vector exists, and give both a linear program and an iterative procedure for finding it. They then proceed with an equilibrium analysis and show the existence of a *complete-information* Nash equilibrium that results in the maximum total payment (this does not imply that the mechanism is incentive compatible). A Web-based system based on their model has been implemented at Yahoo! (but it currently requires an internal login).

The Ghosh and Mahdian paper, in turn, is closely related to another paper by us on markets for general settings with externalities [10]. This paper preceded the Ghosh and Mahdian paper, though they were not aware of it when they wrote their paper. In it, we consider general settings where each agent controls a number of variables, and these variables affect the agent's own utility as well as the utilities of others. For example, the variables can represent how much each bidder donates to each charity. Alternatively, they can represent, for example, levels of pollutants emitted by the agents. The goal is to agree on an outcome that is good for all. We study how computationally hard it is to find optimal solutions in this framework, for various definitions of optimality and various restrictions of the setting. The Ghosh and Mahdian setting corresponds to one of the studied versions, where each agent controls one variable (how much the agent donates), externalities are negative (setting the variable selfishly corresponds to a low donation, which negatively affects the other agents), and the objective is to maximize concessions (equivalently, to maximize donations). In this context, we prove the same result about the existence of a unique maximal solution, and give the same iterative algorithm that converges to this solution. We do not study any equilibrium aspects in that paper, though.

12 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-hard 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 valuation functions and payment willingness functions 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. We also considered the case where each charity has a target amount, and bidders' willingness-to-pay functions are concave. Here, we showed that the optimal surplus can be approximated to a ratio m, the number of charities, in polynomial time (and no significantly better approximation is possible in polynomial time unless P=NP); no polynomial-time approximation ratio is possible for maximizing the total donated, unless P=NP. Subsequently, we showed that the clearing problem is much easier when bids are quasilinear (where payment willingness functions are linear)—for maximizing surplus, the problem decomposes across charities, and for maximizing the total donated, a greedy approach is optimal if the bids are concave (although this latter problem is weakly NP-hard when the bids are not concave). For the quasilinear setting, we studied the mechanism design question of incentivizing the bidders to report their preferences truthfully. 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 can be benefits to linking the charities from a mechanism design standpoint. Finally, we discussed an experiment in which we used this methodology to collect money for victims of the 2004 Indian Ocean Tsunami.

There are many directions for future research. One direction is to create a framework that simultaneously generalizes both our framework and the Ghosh and Mahdian framework, so that donations can be conditional on who donates how much to which charities. Another direction is to build a publicly available Web-based implementation of one of these markets. One can also study the scalability of our MIP/LP approach. It may be helpful to identify other classes of bids for which the clearing problem is tractable. Much work remains to be done on the mechanism design problem. Finally, are there good iterative markets for charitable donations?[16]

References

[1] Krzysztof Apt, Vincent Conitzer, Mingyu Guo, and Evangelos Markakis. Welfare undominated Groves mechanisms. In *Proceedings of the Fourth Workshop on Internet and Network Economics (WINE)*, pages 426–437, Shanghai, China, 2008.

[2] Kenneth Arrow. The property rights doctrine and demand revelation under incomplete information. In M Boskin, editor, *Economics and human welfare*. New York Academic Press, 1979.

[3] Lawrence M. Ausubel and Paul Milgrom. The lovely but lonely Vickrey auction. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 17–40. MIT Press, 2006.

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

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

[6] Ruggiero Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 220–229, Chicago, IL, USA, 2008.

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

[8] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 103–110, Edmonton, Canada, 2002.

[9] Vincent Conitzer and Tuomas Sandholm. Applications of automated mechanism design. In *UAI-03 workshop on Bayesian Modeling Applications*, Acapulco, Mexico, 2003.

[10] Vincent Conitzer and Tuomas Sandholm. Expressive negotiation in settings with externalities. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 255–260, Pittsburgh, PA, USA, 2005.

[16] Compare, for example, iterative combinatorial auctions [26, 33].

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

[12] Peter Cramton, Yoav Shoham, and Richard Steinberg. *Combinatorial Auctions*. MIT Press, 2006.

[13] Claude d'Aspremont and Louis-Andre Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.

[14] Michael Garey, David Johnson, and Larry Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1:237–267, 1976.

[15] Arpita Ghosh and Mohammad Mahdian. Charity auctions on social networks. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1019–1028, 2008.

[16] Andrew Gilpin and Tuomas Sandholm. Arbitrage in combinatorial exchanges. In *Agent-Mediated Electronic Commerce (AMEC) workshop*, pages 43–56, New York, NY, 2004.

[17] Devorah Goldburg and Shannon McElligott. Red cross statement on official donation locations. 2001. Press release, http://www.redcross.org/press/disaster/ds_pr/011017legitdonors.html.

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

[19] Mingyu Guo and Vincent Conitzer. Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69–98, 2009. Special Section of Games and Economic Behavior Dedicated to the 8th ACM Conference on Electronic Commerce.

[20] Mingyu Guo and Vincent Conitzer. Optimal-in-expectation redistribution mechanisms. *Artificial Intelligence*, 174(5-6):363–381, 2010.

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

[22] Andreu Mas-Colell, Michael Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[23] Hervé Moulin. Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119, 2009.

[24] Roger Myerson and Mark Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.

[25] George Nemhauser and Laurence Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, 1999. Section 4, page 11.

[26] David Parkes. Iterative combinatorial auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 2, pages 41–77. MIT Press, 2006.

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

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

[29] Michael Rothkopf, Thomas Teisberg, and Edward Kahn. Why are Vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.

[30] Tuomas Sandholm. Issues in computational Vickrey auctions. *International Journal of Electronic Commerce*, 4(3):107–129, 2000. Special Issue on Applying Intelligent Agents for Electronic Commerce.

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

[32] Tuomas Sandholm. Expressive commerce and its application to sourcing: How we conducted $35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–58, 2007.

[33] Tuomas Sandholm and Craig Boutilier. Preference elicitation in combinatorial auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 10, pages 233–263. MIT Press, 2006.

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

[35] John Tagliabue. Global AIDS Funds Is Given Attention, but Not Money. *The New York Times*, 2003. Reprinted on http://www.healthgap.org/press_releases/a03/060103_NYT_HGAP_G8_fund.html.

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

[37] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(1):103–128, 2007.

A Comparison to 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* or the *winner determination 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 are a popular research topic; for an overview, see a recent book summarizing the state of the art [12].

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 they allow for *expressive bidding*. A bidder can express exactly how much each possible *allocation* is worth to her, and thus the globally optimal allocation can be chosen by the auctioneer. Compare this to a bidder having to bid on two different items in two different (single-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 an 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; 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.[17]

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 an expressive charity donation market (as is demonstrated in the main body of the paper).

B Avoiding indirect payments

In an initial implementation, the approach of having donations made out to a central entity (the *center*), and having the center forward these payments to the charities, may not be desirable. Rather, it may be preferable to have a *partially decentralized* solution, where the bidders write out checks to the charities directly according to a solution based on the bids that is computed by the center. In this scenario, the center merely has to verify that bidders 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 bidders 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.

[17]This does not mean that expressive bidding is always *sufficient* for economic efficiency: for example, even when expressive bidding is possible, bidders may strategically misreport their preferences, resulting in economic inefficiency with respect to their true preferences.

To bring this about, we can still use the approach described in the main body of the paper. After we clear the market in the manner described there, we know the amount that each bidder 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 solution. 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. This is likely to lead to complaints and noncompliance with the solution. We can address this issue by letting each bidder specify explicitly (before the clearing of the market) 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).

We can also integrate the direct-payment model into the MIP introduced earlier. To do so, 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, $\pi_{c_i} = \sum_{b_j} \pi_{c_i,b_j}$; and for each b_j, $\pi_{b_j} = \sum_{c_i} \pi_{c_i,b_j}$. 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 formulated in terms of the π_{c_i} and π_{b_j}, as before.

We note that the main part of this paper corresponds to the special case where there are no constraints on which bidders can donate to which charities, so all of the hardness results from the main part of the paper still apply to the direct-payment model.