Truthful Risk-Managed Combinatorial Auctions

Alan Holland and **Barry O'Sullivan**
Cork Constraint Computation Centre
University College Cork, Ireland
{a.holland|b.osullivan}@4c.ucc.ie

Abstract

Given a winning-bid withdrawal in a combinatorial auction, finding an alternative repair solution of adequate revenue without causing undue disturbance to the remaining winning bids in the original solution may be difficult or even impossible. This "bid-takers exposure problem" may be preemptively addressed by finding a solution that is *robust* to winning-bid withdrawal. We introduce the concept of monotonicity-in-expectation. We provide impossibility results concerning truthful mechanisms for robust solutions with bounded social-welfare losses in which the bid-taker cannot rescind items from winning bidders to repair a solution. We also show that this result extends to combinatorial auctions that include a form of leveled-commitment contract. However, we present a positive result regarding truthfulness for combinatorial auctions in a restricted setting that comprises a computationally efficient allocation algorithm that seeks to maximize expected social welfare.

1 Introduction

A combinatorial auction (CA) provides an efficient means of allocating multiple distinguishable items amongst bidders whose perceived valuations for combinations of items differ. Such auctions are gaining in popularity and there is a proliferation in their usage across a variety of industries. Revenue is the most obvious optimization criterion for such auctions, but another desirable attribute is *solution robustness* [Holland and O'Sullivan, 2005a]. A robust solution to a CA is one that can withstand winning-bid withdrawal (a *break*) by making changes easily to form a repair solution of adequate revenue. A *brittle solution* is one in which an unacceptable loss in revenue is unavoidable if a winning bid is withdrawn. Robustness is a preventative measure that protects against future uncertainty by trading-off revenue for solution stability.

The weighted super solutions (WSS) framework has been proposed as a means of finding such solutions [Holland and O'Sullivan, 2005a]. A natural question that arises from that work is whether it is possible to design truthful mechanisms for combinatorial auctions whose social welfare can be risk-managed. It is possible to incentivize truthful bidding in auctions by carefully designing the allocation and payment mechanism so that they satisfy certain properties. The Revelation Principle implies that in a wide variety of settings, only "truthful revelation mechanisms" in which agents truthfully announce their valuations need to be considered when the maximization of a social objective is required. In other words, there are no manipulable mechanisms in which agents strategically report their types that achieve superior outcomes than any non-manipulable mechanism. This principle allows us to focus our attention on truthful mechanisms only. There has been considerable research effort in this field by the computer science community in recent years because of the computational challenges posed by incentive compatibility [Nisan and Ronen, 2001; 2000; Lehmann *et al.*, 2002].

The contributions of this paper are twofold. Firstly, we present impossibility results related to the non-existence of mechanisms to incentivize truthful bidding when robust allocations are required in a combinatorial auction. These results hold for both irrevocable commitments and a form of leveled commitment contract between the bid-taker and bidders. Secondly, we propose an alternative approach to managing risk in combinatorial auctions in which we maximize expected social welfare. We present a positive result regarding truthfulness for CAs in a restricted setting that features a computationally efficient allocation algorithm.

2 Preliminaries

Combinatorial auctions involve a single bid-taker allocating multiple distinguishable items amongst a set of bidders. The bid-taker has a set of m items for sale, $M = \{1, \ldots, m\}$, and bidders submit a set of bids, $\mathcal{B} = \{B_1, \ldots, B_n\}$. A bid is a tuple $B_j = \langle b_j, S_j \rangle$ where $S_j \subseteq M$ is a subset of the items for sale and $b_j \geq 0$ is a bid amount. The winner determination problem (WDP) for a CA is to label all bids as either winning or losing so as to maximize the revenue from winning bids without allocating any item to more than one bid, and is formulated as follows:

$$max \sum_{j=1}^{n} b_j x_j \ s.t. \sum_{j|i \in S_j} x_j \leq 1, \ \forall i \in \{1 \ldots m\}, \ x_j \in \{0,1\}.$$

This problem is $\mathcal{NP}$-complete [Rothkopf *et al.*, 1998] and inapproximable [Sandholm *et al.*, 2005], and is otherwise known as the Weighted Set Packing Problem. The formulation above assumes the notion of *free disposal*. This means that the optimal solution need not necessarily sell all of the items. If the auction rules stipulate that all items *must* be sold, the problem becomes a Set Partition Problem. The WDP has been extensively studied in recent years. The fastest search algorithms that find optimal solutions (*e.g.* CABOB [Sandholm *et al.*, 2005]) can, in practice, solve very large problems involving thousands of bids very quickly. Given the possibility of a winning bid being withdrawn by the associated bidder before the transaction has been completed, large losses in revenue are possible [Holland, 2005; Holland and O'Sullivan, 2005a]. Therefore, it is often important to preempt bid withdrawal by finding a *robust* solution.

DEFINITION 1 (Robust Solution for a CA) *A robust solution for a combinatorial auction is one in which any subset of successful bids whose probability of withdrawal is at least α can be repaired by reallocating items at a cost of at most β amongst other bidders, in order to form a repair solution whose revenue is at least a fraction, γ, of optimal revenue.*

DEFINITION 2 (1-robust Solution for a CA) *A 1-robust solution for a combinatorial auction is one in which any* single *bid whose probability of withdrawal is at least α can be repaired by reallocating items at a cost of at most β amongst other bidders, in order to form a* repair solution *whose revenue is at least a fraction, γ, of optimal revenue [Holland, 2005]. (The definition of a k-robust solution follows in an obvious manner.)*

A solution is considered robust if and only if the revenue of a repair solution is is guaranteed to be at least γ times the optimum for every likely set of winning-bid withdrawals, whose probability of failure is greater than or equal to α. The maximum cost of repair (β), in the case of CAs, reflects the size of any fund for compensatory payments paid to/from the bid-taker following the reassignment of items in order to find a repair solution with satisfactory revenue. If item revocation from bidders by the bid-taker is forbidden, then $\beta = 0$ and compensatory payments to bidders are effectively ∞. Later, we shall consider mutual bid bonds in which a reneging bidder forfeits a bond that is used by the bid-taker to fund compensation to other non-reneging winning bidders when revoking items to form an alternative repair solution [Holland and O'Sullivan, 2005a].

As an example, consider three bids for two items $\langle 100, \{A, B\}\rangle$, $\langle 60, \{A\}\rangle$ and $\langle 60, \{B\}\rangle$. The independent probabilities of each bid being withdrawn if successful are 0.1. If we set $\alpha = 0.05$, $\beta = 0$ and $\gamma = 0.75$, so that all bids are brittle, no non-reneging winning bidder can have an item revoked, and we wish to ensure that a repair solution exists with a minimum revenue of 0.75×120. Then $\langle 1, 0, 0\rangle$ is the optimal 1-robust solution[1] because it can be repaired to form $\langle 0, 1, 1\rangle$, should the winning bid be withdrawn. Note that $\langle 0, 1, 1\rangle$ is not a robust solution because it results in a repair revenue of 60, which is below our threshold, should either of the winning bids be withdrawn.

Most prior literature on the subject of combinatorial auctions assumes that the probability of a winning bid withdrawing from the auction is zero. Therefore, we contend that relaxing this assumption so that the conditional probability of withdrawal in a repair solution is negligible, given that a withdrawal has already occurred in the initial allocation, is a reasonable assumption. We can choose α to establish an upper bound on the number of likely withdrawals that require repairs. A constraint programming approach, called the weighted super solutions framework, for finding such robust solutions has been developed [Holland and O'Sullivan, 2005a; 2005b]. This framework enables us to find such solutions in a computationally feasible manner by facilitating satisfiability tests for repair solutions.

The design of truthful mechanisms for a computationally hard problem such as winner determination in a CA is made difficult by the fact that efficiently computable heuristics cannot be employed whilst still maintaining their non-manipulability characteristics. For example, the class of Vickrey-Clarke-Groves (VCG) mechanisms guarantee that each bidder's dominant strategy is to tell the truth, but it requires solving $q + 1$ optimization problems, where the overall optimal solutions involves q winning bidders. However, non-optimal solutions compromise truthfulness [Nisan and Ronen, 2001]. In this classic setting a necessary condition for any truthful mechanism is that its allocation scheme be monotone.

DEFINITION 3 (Monotone allocation algorithm) *An allocation algorithm is monotone iff whenever a bid v_j wins, an increase in the bid to v'_j still wins, assuming all other bids are fixed.*

3 Non-monotonicity Results

In the case of truthful auctions, however, the payments do not generally equal the declared valuations. For this reason, risk management for such auctions involves managing potential losses in social welfare rather than revenue. Henceforth, when we refer to solution robustness involving truthful bidding we are referring to robustness in terms of *social welfare*. Given exogenous probabilities of failure, we adopt a concept of monotonicity-in-expectation to analyze allocations that support contingency planning. Clearly, in our setting a necessary condition for any truthful mechanism is that its allocation scheme be monotone-in-expectation.

DEFINITION 4 (Monotone-in-expectation allocation algorithm) *An allocation algorithm is monotone-in-expectation if whenever a bid v_j wins with exogenous probability p, an increase in the bid to v'_j means that it wins with at least probability p, assuming all other bids are fixed.*

3.1 Robust-Allocation Algorithms

We show that an increase in a bid amount may provide a repair solution for a previously irreparable bid. This enlarges the set of bids that may be considered when determining an optimal robust solution. We shall later see how this can occur in practice in Example 1. Lemma 1 provides an impossibility

[1]The notation $\langle 1, 0, 0\rangle$ means that only the first bids wins.

result regarding the monotonicity-in-expectation of robust-allocation algorithms when finding an optimal 1-robust solution. The non-existence of a monotone-in-expectation allocation algorithm implies that for some bidder a, if v_a is a winning declaration with some probability, p, then for some higher declaration $v'_a > v_a$ it is more likely to become a losing declaration. The proof utilizes the fact that if v_a becomes a losing bid with a higher probability, then it must form part of a repair allocation for another previously irreparable bid that partakes in the subsequent winning allocation.

LEMMA 1 *Any allocation algorithm for finding a social-welfare maximising 1-robust solution to a CA with irrevocable assignments of items is non-monotone-in-expectation.*

Proof: Example 1 suffices to prove this result. ■

EXAMPLE 1 Consider bids $v_1\langle 13, A\rangle$, $v_2\langle 12, B\rangle$, $v_3\langle 8, C\rangle$, $v_4\langle 33, AB\rangle$ and $v_5\langle 24, BC\rangle$ (Table 1). Let $\alpha = 0.1$ and all bids are robust, *i.e.* not likely to be withdrawn, except for v_4 whose probability of withdrawal, 0.15, exceeds the threshold value of α. In this example we consider irrevocable assignments so we have no budget to revoke items from non-reneging winning bidders and compensate them for such actions, so $\beta = 0$. We let $\gamma = \frac{34}{41}$, so the minimum acceptable social welfare is 34. The optimal robust solution is, therefore, $\{v_1, v_5\}$ with social welfare of 37, because $\{v_3, v_4\}$, the solution with the highest aggregate declared value to bidders, cannot be repaired to form a solution of at least 34 should v_4 be withdrawn. The optimal repair solution, $\{v_1, v_2, v_3\}$, attains only 33.

Table 1: Combinatorial auction for Example 1.

	Item Combinations					
Bids	A	B	C	AB	BC	Withdrawal prob
v_1	13	0	0	0	0	$0.02(< \alpha)$
v_2	0	12	0	0	0	$0.03(< \alpha)$
v_3	0	0	8	0	0	$0.04(< \alpha)$
v_4	0	0	0	33	0	$0.15(\geq \alpha)$
v_5	0	0	0	0	24	$0.05(< \alpha)$

Let us suppose that v_1 becomes $v'_1\langle 16, A\rangle$, so that the declaration has increased by 3. A monotone-in-expectation allocation algorithm requires that v_1 remains a winning bid in the initial allocation because if it were to become part of a repair allocation the bidder would have a lower probability of receiving the item. However, allocation $\{v_3, v_4\}$ can now be repaired to form a solution with a social welfare of 36 should v_4 be withdrawn. Therefore, the new optimal 1-robust solution becomes $\{v_3, v_4\}$. This implies that v'_1 is less likely to be awarded item A following the increase from v_1, violating the monotonicity-in-expectation requirement. △

Lemma 1 demonstrates that increasing a winning bid has the possible side-effect of creating new robust solutions by forming a repair solution for a previously irreparable allocation of higher social welfare. Example 1 shows how such allocation algorithms can cause the probability of a declaration's success to decrease following an increase in its bid amount.

THEOREM 1 *A normalized[2] truthful mechanism is impossible for 1-robust solutions where item assignments are irrevocable.*

Proof: A normalized mechanism is truthful if and only if its allocation algorithm is monotone and its payment scheme is based on the critical value [Mu'alem and Nisan, 2002]. When considering exogenous probabilities of contingency allocations the algorithm must be monotone-in-expectation. Lemma 1 shows that a monotone-in-expectation allocation algorithm is impossible to achieve for 1-robust solutions where item assignments are irrevocable. ■

Theorem 1 shows that it is impossible to incentivize truthful bidding to achieve 1-robust solutions. Monotonicity-in-expectation is hindered by the fact that irreparable bids are discounted from consideration when determining the winners.

COROLLARY 1 *A normalized truthful mechanism is impossible for k-robust solutions where item assignments are irrevocable.*

In prior work, the lower bound on tolerable revenue was introduced to permit satisfiability testing for repairs in order to improve computational feasibility [Holland and O'Sullivan, 2005a]. However, this constraint inhibits truthful mechanisms because an increase in a bid amount for a bid participating in an optimal robust solution can simultaneously cause a brittle allocation of higher revenue to become robust. We need to circumvent this problem in order to find a truthful mechanism for risk-managed solutions as opposed to robust solutions that provide *guarantees* on the social welfare of repair allocations.

The WEIGHTED-SUPER-SOLVE algorithm [Holland and O'Sullivan, 2005b] is an example of an allocation algorithm that finds such robust solutions. Although the allocation problem is $\mathcal{NP}$-complete when a fixed bound is placed on the size of any break, its application within a VCG-based mechanism would remain of the same asymptotic complexity. An obvious, yet interesting, question remained as to whether it could induce truthful bidding. Theorem 1 informs us that such an algorithm is non-monotone-in-expectation so the answer to this question is negative.

3.2 Mutual Bid Bonds

Some auction solutions are inherently brittle and it may be impossible to find a robust solution. We can alter the rules of an auction so that the bid-taker can revoke items from winning bidders, to be reallocated so that the reparability of solutions to such auctions can be improved [Holland and O'Sullivan, 2005a]. In this section we investigate whether it is possible to allocate items in a robust solution using an auction model that permits bid and item withdrawal by the bidders and bid-taker, respectively, whilst incentivizing truthful bidding.

We adopt a model that incorporates *mutual bid bonds* to enable solution reparability. Mutual bid bonds are a form of insurance against the winner's curse for the bidder whilst also compensating bidders in the case of item withdrawal from

[2] A *normalized* mechanism ensures that non-winners pay zero.

winning bids. Such "*Winner's Curse & Bid-taker's Exposure*" insurance comprises a fixed, non-zero fraction, κ, of the bid amount for all bids. Mutual bid bonds are mandatory for each bid in our model. The conditions attached to the bid bonds are that the bid-taker be allowed to annul winning bids (item withdrawal) when repairing breaks elsewhere in the solution. In the interests of fairness, compensation is paid to bidders from whom items are withdrawn and is equivalent to the penalty that would have been imposed on the bidder should he have withdrawn the bid.

If the decommitment penalties are the same for both parties in all bids, κ does not influence reparability. It merely influences the levels of penalties and compensation transacted by the agents. Low values of κ incur low bid withdrawal penalties and simulate a dictatorial bid-taker who does not adequately compensate bidders for item withdrawal. Increased levels of winning-bid withdrawal are likely when the penalties are low in a common or affiliated-values model. High values of κ tend towards full-commitment. The penalties paid are used to fund a reassignment of items to form a repair solution of sufficient value by compensating previously successful bidders for withdrawal of items from them.

LEMMA 2 *Robust-allocation algorithms for finding an optimal 1-robust solution to a CA with mutual bid bonds are non-monotone-in-expectation.*

Proof: Example 2 suffices to prove this result. ■

We present an example to demonstrate how an increase in a declaration can result in its likelihood of success being decreased because it forms part of a repair solution for a previously irreparable solution. Note that Example 2 has different bid amounts for v_3 and v_5 than those in Example 1. These were altered so that the withdrawal of v_4 would provide insufficient funds to repair $\langle v_3, v_4\rangle$ by revoking items from v_3. Mutual bid bonds only permit revocation of items from a set of winning bids whose summation of declarations is less than that of the withdrawn bid.

EXAMPLE 2 Consider bids $v_1\langle 13, A\rangle$, $v_2\langle 12, B\rangle$, $v_3\langle 38, C\rangle$, $v_4\langle 33, AB\rangle$ and $v_5\langle 54, BC\rangle$ (Table 2). Again, we let $\alpha = 0.10$ so that all bids are robust except for v_4 whose probability of withdrawal is 0.12. We let $\gamma = \frac{64}{71}$, so the minimum acceptable social welfare for a solution is 64. The bid-taker would like to be able to award the items to $\{v_3, v_4\}$ but v_4 is a brittle bid that cannot be repaired satisfactorily. The optimal robust solution is, therefore, $\{v_1, v_5\}$ with welfare of 67.

Table 2: Combinatorial auction for Example 2.

	Item Combinations					
Bids	A	B	C	AB	BC	Withdrawal prob
v_1	13	0	0	0	0	$0.01(<\alpha)$
v_2	0	12	0	0	0	$0.02(<\alpha)$
v_3	0	0	38	0	0	$0.03(<\alpha)$
v_4	0	0	0	33	0	$0.12(\geq\alpha)$
v_5	0	0	0	0	54	$0.04(<\alpha)$

Suppose that v_1 becomes $v_1'\langle 16, A\rangle$, so the declaration has increased by 3. The mutual bid bond provides a fund of $\beta = \kappa \times 33$ if v_4 is withdrawn, but this is insufficient to allow revocation of items from v_3 who requires a compensatory payment of $\kappa \times 38$. A monotone-in-expectation allocation algorithm requires that v_1' has least the same probability of receiving item A as it did previously (Definition 3). However, v_4 can now be repaired to form a solution, $\langle v_1, v_2, v_3\rangle$, of welfare 66 should the bid be withdrawn, therefore, the new optimal 1-robust solution following the increase in v_1 is $\{v_3, v_4\}$. This implies that v_1' becomes a losing declaration with a greater probability following the increase from v_1, thus violating the monotonicity-in-expectation requirement. △

THEOREM 2 *A normalized truthful mechanism is impossible for 1-robust solutions and mutual bid bonds.*

Proof: Lemma 2 shows that a monotone-in-expectation allocation algorithm is impossible to achieve for 1-robust solutions with mutual bid bonds. ■

COROLLARY 2 *A normalized truthful mechanism is impossible for k-robust solutions where items can be revoked based on mutual bid bonds.*

4 Truthful Risk-Managed Allocations

In this section we investigate ways in which we can compromise our allocation algorithm and restrict the CA setting so that we can achieve a truthful mechanism. Various polynomial-time approximation algorithms can provide good or near optimal solutions very quickly. However, Nisan and Ronen [2000] showed that a non-optimal solution can in fact be *degenerate* so that results can be arbitrarily far from the optimum, in terms of social efficiency. A seminal positive result, due to Lehmann *et al.* [2002], showed that by restricting the set of agents' preferences to be what they termed, *single-minded*, i.e. agents are only interested in a single bundle of items, it is possible to develop greedy mechanisms that are both truthful and computationally efficient.

DEFINITION 5 (Single-minded bidder) *Bidder j is single-minded if there is a set of goods $S_j \subseteq M$ and a value $v_j^* > 0$ such that* $v_j(S) = \begin{cases} v_j^* & S \supseteq S_j \\ 0 & otherwise. \end{cases}$

Lehmann *et al.* showed, axiomatically, that a mechanism is truthful, given single-minded bidders, if it fulfils the following requirements: the mechanism's allocation algorithm is monotone (Definition 3); the bidder is either awarded his set of desired items and no more or nothing at all (*exactness*); a winning bidder pays the lowest value he could have declared to win the items (*critical value*); a bidder never pays less for a superset of items in his bid (*payment monotonicity*); losing bidders pay zero (*participation constraints*).

We adopt a similar model to Mu'alem and Nisan [2002] to investigate possible mechanisms that incentivize truthful bidding over a set U of m items. Each bidder j, of n bidders in total, has a non-negative valuation function v_j for a subset $S_j \subseteq U$ of items. We assume that bidders' valuation functions are private, as is commonplace within mechanism design [Archer and Éva Tardos, 2001] and probabilities of withdrawal are exogenous.

Informally, single-minded bidders are willing to pay v_j^*, a privately known valuation, on condition that they minimally receive a set of items, S_j. This condition is necessary to satisfy the individual rationality constraints. Mu'alem and Nisan [2002] introduced *known* single-minded bidders, whereby the subsets S_j are known to the mechanism. In fact, only the cardinality of these subsets needs to be known to the mechanism. This mechanism is composed of an allocation algorithm, $A(v)$, whose inputs are the bid declarations $v = \{v_1, \ldots, v_n\}$, the respective desired items $\{S_1, \ldots, S_n\}$ and a payment rule $p(v)$. The output of $A(v)$ is a subset of pair-wise disjoint winning bids.

We present a positive result: a *truthful approximation scheme whose allocation scheme accounts for the risk of potential single winning-bid withdrawals.* We develop an approximation algorithm similar to the LP-based allocation algorithm of Mu'alem and Nisan [2002] and show that it is monotone-in-expectation. The IP formulation for maximizing expected social welfare, given the possibility of single bid-withdrawal, is the following:

$$\max \left(\sum_{j=1}^{n} (1-p_j) v_j x_j + \sum_{k=1}^{n} \sum_{j=1}^{n} p_k v_j r_{kj} \right) \tag{1}$$

where $r_{kj(k \neq j)} \in \{0,1\}, x_j \in \{0,1\}, r_{kk} \in \{0\}$ and is subject to the following constraints:

$$\sum_{j|i \in S_j} x_j \leq 1 \quad \forall i \in \{1,..,m\}$$

$$\sum_{j|i \in S_j, j \neq k} x_j + r_{kj} \leq 1 \quad \forall i \in \{1,..,m\}, k \in \{1,..,n\}$$

where $0 \leq p_j < \frac{1}{2}$ is the probability of bid j being withdrawn so that we do not expect the repair allocation to have higher revenue than the initial allocation. x_j is the decision variable for bid j, v_j the amount of bid j and r_{kj} is 1 iff bidder j receives her bundle if bidder k withdraws but does *not* receive it if bidder k does not withdraw. We assume conditional probabilities of failure so that the likelihood of multiple winning-bid withdrawals is zero. Notwithstanding this assumption, the formulation still remains computationally infeasible because of the size of the input, which requires $\mathcal{O}(n^2)$ variables. A branch and bound search to find an optimal solution quickly becomes prohibitively expensive as n increases. The LP relaxation for Equation 1 relaxes the integrality constraints on x_j and r_{kj} so that bids, in effect, can partially win. Although this problem also has $\mathcal{O}(n^2)$ input variables, it can be solved in polynomial time. The LP formulation, denoted $LP(v)$, is the same as the integer formulation except that we denote the decision variables as y_j (instead of x_j) and y_j, r_{kj} and r_{kk} are $\in [0,1]$. Note that prior work on robust optimization concerns uncertainty about the precision in coefficients in an ILP [Bertsimas and Sim, 2004], and is therefore not appropriate for our problem that concerns uncertainty about the validity of assignments to variables.

There are no constraints on minimum social welfare in this model. These constraints were originally introduced as a compromise on the original problem of utility maximization so that we could improve computational feasibility whilst still maintaining the integrality constraints on winning bids [Holland and O'Sullivan, 2005a]. By removing these lower bounds on repair revenues and optimizing expected utility using the LP relaxation, we need to then determine the allocation from the results of this non-integer solution. It is natural to assume that y_j, representing the status of bid v_j in the optimal allocation of the linear relaxation, would provide a good heuristic for guiding the choice of an approximate solution because it reflects a fraction of how much that bid should win in an optimal solution. In general, however, such heuristics are not necessarily truthful.

Algorithm 1: LP-RISK-MANAGE

```
input  : Bids V
output: Allocation ⟨X, R⟩
begin
    Y ← LP(V)  // Optimal LP solution
    foreach j ∈ {1...n} do
        if y_j(1 − p_j) + Σ_{k=1}^{n} r_kj p_k > 1/2 then
            X[j] ← 1
    foreach j, k ∈ {1...n} do
        R[k][j] ← 0
end
```

Informally, LP-RISK-MANAGE (Algorithm 1) assigns any bid whose value in the corresponding LP solution is *expected* to be greater than $\frac{1}{2}$, when considering the exogenous probabilities of withdrawals, to be a winning bid. In order to prove the monotonicity-in-expectation of Algorithm 1 it is necessary to show that for any v_{-j}, where $v_{-j} = v \setminus \{v_j\}$, y_j is a non-decreasing function of v_j.

LEMMA 3 *$y_j(1-p_j) + \sum_{k=1}^{n} r_{kj} p_k$ is a non-decreasing function of v_j for any fixed v_{-j}.*

First let us define $U(x,r,v) = \sum_{j=1}^{n}(1-p_j)v_j x_j + \sum_{k=1}^{n}\sum_{j=1}^{n} p_k v_j r_{kj}$ for the purposes of notational brevity.

Proof: Consider optimal feasible solutions $\mathcal{Y}$ and $\mathcal{Y}'$ to the linear program, $LP(v)$ and $LP(v')$, respectively, where $v_j \leq v_j'$ and $\Delta = v_j' - v_j$. $\mathcal{Y}'$ is a feasible, though non-optimal, solution to $LP(v)$ and so $U(y', r', v) \leq U(y, r, v)$. $\mathcal{Y}$ is also a feasible solution to $LP(v')$ and so $(1-p_j)\Delta y_j + \sum_{k=1}^{n} p_k \Delta r_{kj} + U(y,r,v) \leq (1-p_j)\Delta y_j' + \sum_{k=1}^{n} p_k \Delta r_{kj}' + U(y', r', v) \therefore 0 \leq U(y,r,v) - U(y',r',v) \leq (y_j' - y_j)(1-p_j)\Delta + \sum_{k=1}^{n}(r_{kj}' - r_{kj})p_k \Delta$. The expression on the very right hand side is ≥ 0 and we can rearrange it so that the (y_j', r_{kj}') and (y_j, r_{kj}) components are on opposite sides, $y_j'(1-p_j) + \sum_{k=1}^{n} r_{kj}' p_k \geq y_j(1-p_j) + \sum_{k=1}^{n} r_{kj} p_k$. ■

LEMMA 4 *Algorithm 1 (*LP-RISK-MANAGE*) is monotone-in-expectation.*

Proof: v_j is a winning declaration if and only if $\frac{1}{2} < y_j(1-p_j) + \sum_{k=1}^{n} r_{kj} p_k$. From Lemma 3 we know that $\frac{1}{2} < y_j(1-p_j) + \sum_{k=1}^{n} r_{kj} p_k \leq y_j'(1-p_j) + \sum_{k=1}^{n} r_{kj}' p_k$ so any $v_j' > v_j$ is also a winning declaration. ■

THEOREM 3 *A normalized mechanism that assigns items to bidders according to Algorithm 1 (*LP-RISK-MANAGE*) is truthful-in-expectation when its payment scheme is based upon the critical value.*

Proof: We know from [Mu'alem and Nisan, 2002] that for any v_{-j} and monotone allocation algorithm there exists a single critical value below which v_j is always a losing declaration and above which it is always winning. A normalized mechanism is truthful when its payment scheme is based upon the critical value, and the allocation algorithm is monotone [Mu'alem and Nisan, 2002]. From Lemma 4 we can conclude that LP-RISK-MANAGE is monotone-in-expectation, thereby weakening the solution concept to truthfulness-in-expectation. ■

The most important application of an approximation algorithm is within a mechanism that can provide guarantees on algorithm performance. Algorithm 2 provides a $(m + 1)$-approximation for expected social-welfare maximizing CAs when single winning-bid withdrawal is possible.

Algorithm 2: $(m + 1)$-APPROX-RISKMANAGE

```
input  : Bids V
output: Allocation ⟨X, R⟩
begin
    //Highest expected bid amount
    j ← arg max(p_j × v_j)
    //Second Highest expected bid amount as a repair
    j_2 ← arg max(p_{j_2} × v_{j_2}, j_2 ≠ j)
    X[a] = 0, R[b][c] = 0   ∀a, b, c = {1..n}
    X[j] = 1, R[j][j_2] = 1
    //σ(LP(V)): the optimal objective value
    if σ(LP(V))/(m+1) < (1 − p_j)v_j + p_j v_{j_2} then
        return ⟨X, R⟩
    else return LP-RISK-MANAGE(V)
end
```

THEOREM 4 *Algorithm 2 provides an* $m + 1$ *approximation of the optimal outcome according to Equation 1.*

Proof: (Sketch): In the worst-case situation for Algorithm 2, the largest single bid amount is $\frac{1}{m}$ times the optimal initial allocation in terms of aggregate declared value. All repair bids for the optimum will result in no loss in declared value, and thus the second highest declaration will equal this amount when there are two or more bids. ■

5 Conclusion

We proved a key negative result using the notion of 1-robust solutions: robust-allocation algorithms are non-monotone-in-expectation for the case in which no items can be revoked from winning bidders. This result extends to k-robust solutions. We also showed that they remain non-monotone-in-expectation when mutual bid bonds, a class of leveled commitment contract, are adopted.

We circumvented the impossibility results for the case in which no items can be revoked by the bid-taker and bidders are known single-minded, by removing the constraints on minimal social welfare. When the conditional probability of multiple withdrawals is set to zero, we used a linear programming-based approximation algorithm to find approximately utility-maximizing solutions. We outlined a computationally efficient $(m + 1)$-approximation scheme that is truthful-in-expectation for risk managed solutions. The weak bound on the approximation ratio is indicative of the difficulty in attaining truthful risk-managed allocations but serves as a useful starting point for future work in this nascent research area.

Acknowledgements

This work has received support from Enterprise Ireland under grant number PC/2005/0117.

References

[Archer and Éva Tardos, 2001] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of FOCS*, pages 482–491, October 2001.

[Bertsimas and Sim, 2004] Dimitris Bertsimas and Melvyn Sim. Robust discrete optimization and network flows. *Mathematical Programming*, 98:43–71, 2004.

[Holland and O'Sullivan, 2005a] Alan Holland and Barry O'Sullivan. Robust solutions for combinatorial auctions. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, pages 183–192, May 2005.

[Holland and O'Sullivan, 2005b] Alan Holland and Barry O'Sullivan. Weighted super solutions for combinatorial auctions. In *AAAI-05*, pages 378–383, July 2005.

[Holland, 2005] Alan Holland. *Risk Management for Combinatorial Auctions*. PhD thesis, University College Cork, Ireland, July 2005.

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

[Mu'alem and Nisan, 2002] Ahuva Mu'alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In *Proceedings of AAAI-02*, pages 379–384, 2002.

[Nisan and Ronen, 2000] Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. In *ACM Conf. on Electronic Commerce*, pages 242–252, 2000.

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

[Rothkopf *et al.*, 1998] Michael H. Rothkopf, Aleksander Pekeč, and Ronald M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.

[Sandholm *et al.*, 2005] Tuomas Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. CABOB: A fast optimal algorithm for winner determination in combinatorial auctions. *Management Science*, 51(3):374–390, 2005.