

A Cooperative Approach to Collusion in Auctions

YORAM BACHRACH, Microsoft Research

and

MORTEZA ZADIMOGHADDAM, MIT

and

PETER KEY, Microsoft Research

The elegant Vickrey Clarke Groves (VCG) mechanism is well-known for the strong properties it offers: dominant truth-revealing strategies, efficiency and weak budget-balance in quite general settings. Despite this, it suffers from several drawbacks, prominently susceptibility to collusion. By jointly setting their bids, colluders may increase their utility by achieving lower prices for their items. The colluders can use monetary transfers to share this utility, but they must reach an agreement regarding their actions. We analyze the agreements that are likely to arise through a cooperative game theoretic approach, transforming the auction setting into a cooperative game. We examine both the setting of a multi-unit auction as well as path procurement auctions.

Categories and Subject Descriptors:

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*;

J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Algorithms, Theory, Economics

Additional Key Words and Phrases: Collusion, Cooperative Game Theory, Bid Rigging

1. INTRODUCTION

Auctions are a commonly used mechanism for selling or allocating goods. A key problem that such mechanisms face is that bidders may not bid truthfully. Such bid-shading is common in many auctions. Under such strategic behavior, even a mechanism that attempts to maximize social welfare based on the information it is given may reach a sub-optimal allocation, as it is given incorrect information. By using a proper payment rule, it is possible to incentivise the bidders to truthfully report their valuations. The most prominent method for achieving this is the VCG mechanism [Vickrey 1961; Clarke 1971; Groves 1973]. A detailed introduction to VCG and its properties is contained in [Nisan et al. 2007].

Despite its advantages, VCG has many shortcomings [Ausubel and Milgrom 2006], including vulnerability to collusion [Milgrom 2004]. *Collusion* is an agreement between two or more agents to limit competition by manipulating or defrauding in order to obtain an unfair advantage. Such manipulations include agreements to divide the market, set prices or limit production or bids. One prominent form of collusion is bid rigging, where participants in an auction change their bids in an attempt to lower prices. Collusion and anti-competitive behavior occur in many domains, and many of its forms are illegal [Milgrom 2004].

Forming a colluding coalition in an auction may be a difficult task, however there are several possible approaches, such as bidding rings and bidding clubs [Leyton-

Authors' address: yorambac@gmail.com zadimoghaddam@gmail.com Peter.Key@microsoft.com

Brown et al. 2002; 2000]. We consider forming a colluder coalition when the colluders can form binding agreements and under the simplifying assumption that the colluders have information regarding the valuations of the other colluders and the other participants. Obviously, in real life situations there is likely to be much uncertainty regarding the valuations, however in some cases a certain subset of the agents may be better informed and may wish to collude in the auction.

We focus on the agreements regarding sharing the gains from collusion, using solution concepts from cooperative game theory. Specifically, we consider collusion in multi-unit auctions [Bachrach 2010] and in path procurement auctions [Bachrach et al. 2010] under the Vickrey-Clarke-Groves (VCG) mechanism.

In a multi-unit domain there are multiple *identical* items to be allocated. Each agent thus only cares about the *number* of items she receives. We assume free disposal, so agents always value obtaining more items at least as much as obtaining less items. Agents can thus express their preferences as a valuation function mapping the number of items they obtain to their valuation for this quantity of items. In path procurement auctions, a buyer procures a path from a source s to a target t in a graph $G = (V, E)$. Each edge $e_i \in E$ is owned by a_i , who incurs a cost c_i when her edge is used. The cost c_i is known only to a_i . The buyer must compensate edges on the chosen path for their costs.

Given the agents' preferences a mechanism can take the optimal decision: in multi-unit auctions, given the valuation functions, the mechanism can use the optimal allocation; in path procurement auctions, given the true edge costs the mechanism can find the path of minimal costs. However, this information is private and known only to the agents and their reports may be viewed as bids in an auction. Agents may bid strategically, misreporting their valuation functions to achieve a better outcome for themselves. The VCG scheme is the canonical method for incentivising the agents to bid truthfully (reveal their true valuation function). Although in VCG auctions any single agent is incentivised to truthfully reveal her valuation function, *collusion* might still be beneficial: *several* agents may agree to misreport their valuations in a *coordinated* way, and split the gains from this manipulation.

We analyze how the colluders might share the gains from such a manipulation using cooperative game theory, modeling the domain as a coalitional game called the *collusion game*. We examine the possible agreements using solutions such as the Shapley value [Shapley 1953] and the core [Gillies 1953]. We show that in multi-unit auctions where the marginal valuations are diminishing the colluders can form *stable* agreements and can even split the gains in a “fair” manner. On the other hand, in path procurement auctions, there are cases where the colluders cannot form long-lasting agreements, providing a barrier against collusion.

2. THE MODEL

Consider a multi-unit auction, where the auctioneer offers to sell t identical items to n bidders, $\{1, \dots, n\}$. Each bidder i has a certain valuation to any number of items she receives, given by a function $v_i : \mathbb{N} \cup \{0\} \rightarrow \mathbb{R}$ which maps the number of units to the utility to that bidder. We denote the marginal value of the j 'th item to agent i as $m_i(j) = v_i(j) - v_i(j - 1)$. We assume free disposal, so for any $k > j$ we have $v_i(k) > v_i(j)$. We say the auction has diminishing marginal valuations if

for any agent i and number of items j , $v_i(j+2) - v_i(j+1) \leq v_i(j+1) - v_i(j)$. Since items are identical, an allocation simply maps an agent to the number of items she receives, so it is a vector of quantities $q = (q_1, \dots, q_n)$ such that $\sum_{i=1}^n q_i = t$. Given an allocation q , we denote the total utility the allocation q generates as $f(q) = \sum_{i=1}^n v_i(q_i)$. The total utility can be expressed as the following sum of marginal utilities: $f(q) = \sum_{i=1}^n \sum_{j=1}^{q_i} m_i(j)$.

Denote the marginals list for agent i as $m_i = (m_i(1), \dots, m_i(t))$. Given functions v_1, \dots, v_n of the bidders, we can compute the optimal allocation of items, maximizing $f(q)$. However, the v_i functions (or the marginal lists m_i) are private information of the bidders. In such settings, the Vickrey-Clarke-Groves (VCG) mechanism can be used to induce truthfulness. VCG chooses the optimal allocation according to the reported valuations, and charges each agent its impact on the social welfare of the others. When the marginal valuations are diminishing, a greedy approach finds the optimal allocation. This method obtains a list of marginal utilities for each of the bidders. It then sorts these lists, from highest to lowest, but keeps track of the agent from which the marginal originated from. When computing an optimal allocation for t items, it takes the first t items in the sorted lists, and assigns items to the agents from which these marginals originated from.

Given two marginal lists, $m_i = (m_i(1), \dots, m_i(t))$ and $m_j = (m_j(1), \dots, m_j(t))$ we denote the concatenated list $(m_i, m_j) = (m_i(1), \dots, m_i(t), m_j(1), \dots, m_j(t))$. Given a marginal list m , we denote the sorted list as m_s . Given two marginal lists m_i, m_j , we can concatenate them and sort the concatenated list to obtain $(m_i, m_j)_s$ and denote $(m_i, m_j)_s = (m_i \cup m_j)_s$. Given a subset C of agents, we denote the sorted set of marginals of all the agents as $m_s^C = (\cup_{i \in C} m_i)_s$ (i.e. the marginal list generated by concatenating all the marginal lists and sorting the resulting list). We denote the sorted list or marginals of all the users *except* user j as $m_s^{-j} = (\cup_{i \in C \setminus \{j\}} m_i)_s$. We denote the sorted marginals of all the agents that are *not* in C as $m_s^{-C} = (\cup_{i \notin C} m_i)_s$. Using this notation we can easily express the VCG prices. Denote sorted list of marginals for all the users except i as $m_s^{-i} = (b_1, b_2, \dots, b_{(n-1)t})$. We note that the optimal allocation for $N \setminus \{i\}$ has a social welfare of $\sum_{i=1}^t b_i$. When agent i is present, the algorithm uses the full marginals list $m_s^N = (a_1, a_2, \dots, a_{nt})$. The VCG payment for i is thus q_i marginals in m_s^{-i} starting from position $t - q_i$.

One possible manipulation is designating a single agent from the coalition of colluders as the proxy agent, and making sure the proxy agent obtains the number of items allocated to the coalition under truthful revelation. All the marginals of the non-proxy colluders are set to 0. Denote the colluders $C = \{1, 2, \dots, r\}$ with marginal lists m_1, \dots, m_r , and the non-colluders $\{r+1, r+2, \dots, n\}$ with marginals m_{r+1}, \dots, m_n . Each marginal list m_i has t values $m_i(1), \dots, m_i(t)$ and we denote the marginal functions reported by the colluders as $m^{C'} = (m'_1, \dots, m'_r)$. The following is the *optimal* manipulation when the valuation functions are diminishing.

- (1) Compute q_C , the number of items given to C 's members under truthful bids.
- (2) Compute m_s^C the sorted marginals list for the entire colluder coalition.
- (3) Designate agent 1 as the proxy agent, and construct her marginal list as the first q_C values of m_s^C and 0 for any marginal beyond that point: for any $i \leq q_C$ have $m_1(i) = m_s^C(i)$, and for any $i > q_C$ have $m_1(i) = 0$.

(4) The marginals for all the non-proxy colluders j are 0: $m_j(l) = 0$ for any $j \neq 1$.

3. THE COLLUSION GAME

Consider a multi-unit auction of t items, with n agents and a subset $C \subseteq N$, who decide to collude. Under truthful revelation, the VCG mechanism results in the allocation $q^t = (q_1^t, \dots, q_n^t)$ ¹, and payments p_1^t, \dots, p_n^t , so the coalition C as a whole obtains a certain utility: $u^t(C) = \sum_{i \in C} u_i = \sum_{i \in C} v_i(q_i^t) - p_i^t$. If the agents in C decide to collude, they can use *optimal manipulation*, changing the allocation to $q^* = (q_1^*, \dots, q_n^*)$ and the payments to p_1^*, \dots, p_n^* . Under the optimal manipulation, the total number of items C receives does not change, i.e. $q^t(C) = q^*(C)$ where $q^t(C) = \sum_{i \in C} q_i^t$ and $q^*(C) = \sum_{i \in C} q_i^*$. However, the total payment made by the coalition drops, i.e. $p^t(C) \geq p^*(C)$ where $p^t(C) = \sum_{i \in C} p_i^t$ and $p^*(C) = \sum_{i \in C} p_i^*$. The proxy agent can then distribute the items similarly to their distribution under the truthful bids, resulting in colluder i obtaining q_i^t items. The proxy agent is the only colluder who makes a non-zero payment to the mechanism. The colluders must then make intra-coalition monetary transfers, which determine their utility. Thus, the coalition, *as a whole*, generates the following utility to its members $u^*(C) = \sum_{i \in C} v_i(q_i) - \sum_{i \in C} p_i^*$. We define a coalitional game, based on this utility.

Definition 3.1 The Collusion Game. Given a VCG multi-unit auction of t items for agents $N = \{1, 2, \dots, n\}$ with marginal functions m_1, \dots, m_n , we define the value $v(C)$ of a coalition $C \subseteq N$ as²: $v(C) = u^*(C)$.

Once the coalitional game is defined, we can focus on determining how the total utility generated by the colluders can be shared amongst the members. Monetary transfers allow any distribution of the utility, regardless of how the items are allocated. Cooperative game theory offers solution concepts, predicting which agreements are likely between the colluders. When the marginal valuation functions are diminishing the above collusion game is convex. Due to the convexity of the game, it always has a *non-empty core* and there is a simple polynomial algorithm that finds core imputations. Also, the Shapley value which is widely considered a “fair” solution, is in the core. Thus, the colluders can share their gains from the manipulation in a stable and “fair” manner³. This makes collusion a significant problem in such auctions. Our results regarding “optimal collusion” are similar to merging in weighted voting games [Bachrach and Elkind 2008; Aziz and Paterson 2009]. This attack operates similarly to the techniques discussed regarding false name proofness and self enforcing collusion [Yokoo et al. 2005; Conitzer and Sandholm 2006].

Although the Shapley value is a fair and stable agreement for the colluders, the colluders must be able to tractably compute it to make it an appealing agreement for them. A multi-unit auction where all of the marginal functions m_i are identical except for a *constant* number b of agents, a_1, \dots, a_b who have a different marginal function is called a b -bounded collusion domain. In a b -bounded collusion domain,

¹The subscript t stands for truthful.

²In this definition v maps a coalition to the utility they achieve, so v denotes the characteristic function, not to be confused with $v_i(q_i)$ which is a valuation of a number of items.

³We mean “fair” for the colluders, of course. Collusion is very *unfair* for the auctioneer.

the Shapley value can be computed in polynomial time, and testing whether a given imputation is in the core can also be performed in polynomial time.

In collusion games, a change in the number of available items can affect the agent's power (or share of the gains) in the game, similarly to the fact that changes to the quota in weighted voting games have also been shown to influence power [Zuckerman et al. 2008]. The optimal collusion scheme allows colluders to compute and form stable collusion agreements due to the nature of the auction.

4. PATH PROCUREMENT AUCTIONS

The VCG mechanism can also be used in path procurement auctions (PPAs) where the buyer procures edges $P \subseteq E$ forming an $s-t$ -path from a set of agents, each owning an edge in the graph. Each agent has a cost c_i that it incurs if its edge is used by the buyer and the mechanism asks each e_i to provide a bid b_i for using the edge. If the agent is truthful, she would report c_i . Given the edges' true costs, one can find the minimal cost $s-t$ -path, but the costs are private information. Again, VCG can be used to induce truthfulness. Let $G = (V, E)$ be a path procurement domain, with cost c_i for edge $e_i \in E$, and let b_i be the bid of e_i . Denote the minimal cost path (according to the declared b_i 's) as $(e_{i_1}, e_{i_2}, \dots, e_{i_x})$ (of x edges), and let the optimal path not including e_i be $e_{j_1}, e_{j_2}, \dots, e_{j_y}$ (of y edges). Under a VCG auction, if e_i is on the chosen path, the payment to e_i 's owner is $p_i = \sum_{s=1}^y b_{j_s} - \sum_{s=1}^x b_{i_s} + b_i$, otherwise $p_i = 0$. We call an edge coalition *simple* if it only contains edges on the optimal (minimal cost) $s-t$ -path or if it only contains edges on a single simple suboptimal $s-t$ -path. The optimal collusion scheme for simple edge coalitions is where all colluders bid a zero cost for using their edges.

It is possible to define the cooperative collusion game for PPAs as well. We examine a subset $C \subseteq N$, who may decide to collude. Under truthful bidding, VCG chooses path $r_1 = (e_{i_1}, e_{i_2}, \dots, e_{i_x})$ and payments p_1^t, \dots, p_n^t ⁴. If the agents in C decide to collude, they can form a coalition and use a collusion scheme. Denote the chosen path under the optimal manipulation as $r^* = (e_1^*, \dots, e_z^*)$ and the payments under the manipulation p_1^*, \dots, p_n^* . The coalition members gain payments, but the members on the chosen path, $C \cap r^*$, also incur the cost of their edges. Thus, the utility of the colluder coalition C is: $u^*(C) = \sum_{i \in C} p_i^* - \sum_{i \in C \cap r^*} c_i$. Using monetary transfers, the coalition's utility can be distributed among its members in any way they choose. The collusion game for PPAs is defined as follows.

Definition 4.1 Path Procurement Collusion Game. Given a VCG PPA, the value $v(C)$ of a coalition $C \subseteq N$ is: $v(C) = u^*(C)$.

When a coalition forming an $s-t$ -cut colluders, it can obtain any desired price from the auctioneer, so its value is unbounded. However, as opposed to the collusion game for multi-unit auctions with diminishing marginal valuations, the collusion game for PPAs is not always convex. Thus, there are some instances with an empty core. When the core is empty, the colluders have no stable way of sharing their gains, so although there are beneficial manipulations, the colluders are less likely to agree on the monetary transfer. This can serve as a barrier against collusion. Although the core of a PPA collusion game may be empty, if the colluder set is

⁴The subscript t stands for truthful.

restricted to simple edge coalitions, then the collusion game is convex. In this case the game has a non-empty core that contains the Shapley value. Moreover, for this case it is also possible to compute the Shapley value in polynomial time.

5. CONCLUSIONS

We have taken a cooperative approach to analyzing collusion in VCG auctions. Our results indicate that in some auctions even when the colluders have full information regarding each other's preferences and regarding the other participants in the auction, they may not always be able to form *stable agreements*, thus blocking collusion in such auctions. However, in multi-unit auctions the colluders always have a stable agreement when the marginal valuations are diminishing. An interesting future research direction is modeling uncertainty the colluders have regarding the valuations and examining other domains where the inability of the colluders to form stable agreements might mitigate the effects of collusion.

REFERENCES

- AUSUBEL, L. M. AND MILGROM, P. 2006. The lovely but lonely vickrey auction. In *Combinatorial Auctions*. Chapter 1, 17–40.
- AZIZ, H. AND PATERSON, M. 2009. False name manipulations in weighted voting games: splitting, merging and annexation. In *AAMAS*. 409–416.
- BACHRACH, Y. 2010. Honor among thieves: collusion in multi-unit auctions. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 617–624.
- BACHRACH, Y. AND ELKIND, E. 2008. Divide and conquer: False-name manipulations in weighted voting games. In *AAMAS*. 975–982.
- BACHRACH, Y., KEY, P., AND ZADIMOGHADDAM, M. 2010. Collusion in VCG Path Procurement Auctions. In *Proceedings of the Sixth Workshop on Internet and Network Economics*.
- CLARKE, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11, 1, 17–33.
- CONITZER, V. AND SANDHOLM, T. 2006. Failures of the VCG mechanism in combinatorial auctions and exchanges. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*. ACM, 521–528.
- GILLIES, D. B. 1953. Some theorems on n-person games. Ph.D. thesis, Princeton University.
- GROVES, T. 1973. Incentives in teams. *Econometrica* 41, 617–663.
- LEYTON-BROWN, K., SHOHAM, Y., AND TENNENHOLTZ, M. 2000. Bidding clubs: institutionalized collusion in auctions. In *Proceedings of the 2nd ACM conference on Electronic commerce*. ACM, 253–259.
- LEYTON-BROWN, K., SHOHAM, Y., AND TENNENHOLTZ, M. 2002. Bidding clubs in first-price auctions. In *Proceedings of the National Conference on Artificial Intelligence*. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 373–378.
- MILGROM, P. 2004. *Putting Auction Theory To Work*. Cambridge University Press.
- NISAN, N., ROUGHGARDEN, T., TARDOS, E., AND VAZIRANI, V. V. 2007. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA.
- SHAPLEY, L. S. 1953. A value for n-person games. *Contrib. to the Theory of Games*, 31–40.
- VICKREY, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16, 1, 8–37.
- YOKOO, M., CONITZER, V., SANDHOLM, T., OHTA, N., AND IWASAKI, A. 2005. Coalitional games in open anonymous environments. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI 2005)*. 509–514.
- ZUCKERMAN, M., FALISZEWSKI, P., BACHRACH, Y., AND ELKIND, E. 2008. Manipulating the quota in weighted voting games. In *Proceedings of the 23rd national conference on Artificial intelligence-Volume 1*. AAAI Press, 215–220.