False-Name-Proofness with Bid Withdrawal*

(Extended Abstract)

Mingyu Guo
Duke University
Department of Computer Science
Durham, NC, USA
mingyu@cs.duke.edu

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

ABSTRACT

We study a more powerful variant of false-name manipulation in Internet auctions: an agent can submit multiple false-name bids, but then, once the allocation and payments have been decided, withdraw some of her false-name identities (have some of her false-name identities refuse to pay). While these withdrawn identities will not obtain the items they won, their initial presence may have been beneficia to the agent's other identities. We defin a mechanism to be *false-name-proof with withdrawal (FNPW)* if the aforementioned manipulation is never beneficial FNPW is a stronger condition than false-name-proofness (FNP).

We discuss the relation between FNP and FNPW in general combinatorial auction settings. We also propose the *maximum marginal value item pricing (MMVIP)* mechanism, which we prove is FNPW. (The full version contains a number of other results.)

Categories and Subject Descriptors

J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*; I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems

General Terms

Economics, Theory

Keywords

Mechanism design, combinatorial auctions, false-name-proof

1. INTRODUCTION

The line of research on preventing manipulation via multiple fictitiou identities in Internet auctions was explicitly framed by the groundbreaking work of Yokoo *et al.* [9]. Extending *strategy-proofness*—the concept of ensuring that it is always in a bidder's best interest to report her valuation function truthfully—the authors defin an auction mechanism to be *false-name-proof* if the mechanism is not only strategy-proof, but also, under this mechanism, an agent cannot benefi from submitting multiple bids under false names (fictitiou identities). The authors also extended the revelation principle [4] to incorporate false-name-proofness. That is (roughly stated), in settings where false-name bids are possible, it is without loss of generality to focus only on false-name-proof mechanisms.

Several false-name-proof mechanisms have been proposed for general combinatorial auction settings (settings where multiple items are for sale at the same time, and agents can express valuation functions for the items that exhibit substitutability and complementarity).These are the Set mechanism [7], the Minimal Bundle (MB) mechanism [7], and the Leveled Division Set (LDS) mechanism [8].[1] Other work on false-name-proofness includes the following. For general combinatorial auction settings, Yokoo [7] and Todo *et al.* [5] characterized the payment rules and the allocation rules of false-name-proof mechanisms, respectively. False-name proofness has also been studied in the context of voting mechanisms [2, 6]. Finally, Conitzer [1] proposed the idea of preventing false-name manipulation by verifying the identities of certain limited subsets of agents.

Focusing primarily on combinatorial auctions, this paper continues the line of research on false-name-proofness by considering an even more powerful variant of false-name manipulation: an agent can submit multiple false-name bids, but then, once the allocation and payments have been decided, withdraw some of her false-name identities (have some of her false-name identities refuse to pay). While these withdrawn identities will not obtain the items they won, their initial presence may have been beneficia to the agent's other identities, as shown in the following example:

Example 1. There are three single-minded agents $1, 2, 3$ and two items A, B. Agent 1 bids 4 on $\{A, B\}$. Agent 2 bids 2 on $\{B\}$. Let us analyze the strategic options for agent 3, who is single-minded on $\{A\}$, with valuation 1. (That is, $\forall S \subseteq \{A, B\}$, agent 3's valuation for S is 1 if and only if $\{A\} \subseteq S$.) The mechanism under consideration is the VCG mechanism.

If agent 3 reports truthfully, then she wins nothing and pays nothing. Her resulting utility equals 0. If agent 3 attempts "traditional" false-name manipulation, that is, submitting multiple false-name bids, and honoring all of them at the end, then her utility is still at most 0: if 3 wins both items with one identity, then she has to pay at least 4 (while her valuation for the items is only 1); if 3 wins both items with two identities (one item for each identity), then the identity winning $\{B\}$ has to pay at least 2; if 3 wins only $\{B\}$ or nothing, then her utility is at most 0; if 3 wins only $\{A\}$ (in which case $\{B\}$ has to be won by agent 2), then 3's winning identity's payment equals the other identities' overall valuation for $\{A, B\}$

*Guo and Conitzer are supported by NSF under award number IIS-0812113, and by an Alfred P. Sloan Research Fellowship.

Cite as: False-Name-Proofness with Bid Withdrawal (Extended Abstract), Mingyu Guo and Vincent Conitzer, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1475-1476

[1] A very recent paper [3] introduces a new mechanism called the ARP mechanism. However, this mechanism requires the additional restriction that agents are single-minded.

(at least 4), minus 2's valuation for $\{B\}$ (which equals 2). That is, in this case, 3 has to pay at least 2. So, overall, 3's utility is at most 0 if she honors all her bids.

However, agent 3 can actually benefi from submitting multiple false-name bids, as long as she can withdraw some of them. For example, 3 can use two identities, $3a$ and $3b$. $3a$ bids 1 on $\{A\}$. $3b$ bids 4 on $\{B\}$. At the end, $3a$ wins $\{A\}$ for free, and $3b$ wins $\{B\}$ for 2. If 3 can withdraw identity $3b$ (*e.g.*, by never checking that e-mail account anymore), never making the payment and never collecting $\{B\}$, then, she has obtained $\{A\}$ for free.

If we wish to guard against manipulations like the above, we need to extend the false-name-proofness condition. We refer to the new condition as *false-name-proofness with withdrawal (FNPW)*. It requires that, regardless of what other agents do, an agent's optimal strategy is to report truthfully using a single identity, even if she has the option to submit multiple false-name bids, and withdraw some of them at the end of the auction.

To our knowledge, this stronger version of false-name-proofness has not previously been considered. Whether it is more or less reasonable than the original version depends on the context. For example, in an auction, it may be possible to require each participant to place the amount of her bid in escrow, which would prevent manipulation based on withdrawal. However, in some auction contexts, such an arrangement would be too unattractive to the bidders; it also reduces the anonymity of bidding. Additionally, if we are in a setting where the payments are not monetary, but rather are in terms of performance of future services, then it is not possible to put the payments in escrow.

In any case, FNPW is a useful conceptual tool for analyzing false-name-proof mechanisms. This paper also contributes to the research on false-name-proofness in the traditional sense. Since FNPW is stronger than FNP, the MMVIP mechanism we propose in this paper should be of interest in the FNP context as well.

2. RELATION BETWEEN FNP AND FNPW

We start with definin both FNP and FNPW.

Definition 1. FNP. A mechanism is FNP if and only if under it, for any agent, truthful reporting using a single identifie is always at least as good as submitting multiple false-name bids (and honoring all of them).

Definition 2. FNPW. A mechanism is FNPW if and only if under it, for any agent, truthful reporting using a single identifie is always at least as good as submitting multiple false-name bids and then withdrawing some of them.

THEOREM 1. *FNPW is equivalent to FNP plus the following condition:* Others' bids do not help (OBDNH)*: An agent's utility for reporting truthfully does not increase if we add another agent.*

Among the three existing FNP mechanisms for general combinatorial auction settings, the Leveled Division Set (LDS) mechanism [8] does not satisfy OBDNH, while the Set mechanism [7] and the Minimal Bundle (MB) mechanism [7] both satisfy OBDNH. That is, LDS is not FNPW in general, while Set and MB are FNPW.

In Yokoo *et al.* [9], it was shown that the following submodularity condition is sufficien for the VCG mechanism to be FNP.

Definition 3. Submodularity [9]. For any set of bidders Y, whose types are drawn from Θ, for any two bundles S_1, S_2, we have $U(S_1, Y) + U(S_2, Y) \geq U(S_1 \cup S_2, Y) + U(S_1 \cap S_2, Y)$. Here, $U(S, Y)$ is define as the total utility of bidders in Y, if we allocate items in S to these bidders efficientl .

Actually, submodularity is also sufficien for the VCG mechanism to be FNPW. Moreover, unlike for FNP, in the case of FNPW, the converse also holds—if VCG is FNPW, and additionally the type space contains the additive valuations (those valuations which value any bundle at the sum of the values of its elements, with no complementarity and no substitutability), then the type space must satisfy the submodularity condition.

3. MAXIMUM MARGINAL VALUE ITEM PRICING MECHANISM

In this section, we introduce a new FNPW mechanism.

Definition 4. Maximum marginal value item pricing mechanism (MMVIP): Under MMVIP, an agent faces a posted price mechanism. We use $v(i, S)$ to denote agent i's valuation for bundle S. We use $p(i, S)$ to denote the price of bundle S offered to agent i. $p(i, S)$ is define as follows:

$$p(i, S) = \textstyle\sum_{s \in S} p(i, \{s\})$$
$$p(i, \{s\}) = \max_{j \neq i} \max_{T} \{v(j, T \cup \{s\}) - v(j, T)\}$$

That is, MMVIP uses *item pricing*, and the price an agent faces for an item is the maximum possible marginal value that any other agent could have for that item, where the maximum is taken over all possible allocations.

4. REFERENCES

[1] V. Conitzer. Limited verificatio of identities to induce false-name-proofness. In *Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 102–111, Belgium, 2007.

[2] V. Conitzer. Anonymity-proof voting rules. In *Proceedings of the Fourth Workshop on Internet and Network Economics (WINE)*, pages 295–306, Shanghai, China, 2008.

[3] A. Iwasaki, V. Conitzer, Y. Omori, Y. Sakurai, T. Todo, M. Guo, and M. Yokoo. Worst-case efficien y ratio in false-name-proof combinatorial auction mechanisms. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Toronto, Canada, 2010.

[4] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[5] T. Todo, A. Iwasaki, M. Yokoo, and Y. Sakurai. Characterizing false-name-proof allocation rules in combinatorial auctions. In *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 265–272, 2009.

[6] L. Wagman and V. Conitzer. Optimal false-name-proof voting rules with costly voting. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 190–195, Chicago, IL, USA, 2008.

[7] M. Yokoo. The characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 733–742, Acapulco, Mexico, 2003.

[8] M. Yokoo, Y. Sakurai, and S. Matsubara. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence*, 130(2):167–181, 2001.

[9] M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.