False-name-proofness with Bid Withdrawal

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 While these withdrawn identities will not obtain the items they won, their initial presence may have been beneficial to the agent's other identities. We define 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 first give a necessary and sufficient condition on the type space for the VCG mechanism to be FNPW. We then characterize both the payment rules and the allocation rules of FNPW mechanisms in general combinatorial auctions. Based on the characterization of the payment rules, we derive a condition that is sufficient for a mechanism to be FNPW.

We also propose the *maximum marginal value item pricing (MMVIP)* mechanism. We show that MMVIP is FNPW and exhibit some of its desirable properties. We then propose an automated mechanism design technique that transforms any feasible mechanism into an FNPW mechanism, and prove some basic properties about this technique. Since FNPW is stronger than FNP, the mechanisms we obtain in this paper are also FNP. Finally, we prove a strict upper bound on the worst-case efficiency ratio of FNPW mechanisms. In the appendix, we give a characterization of FNP(W) social choice rules.

1. INTRODUCTION

With the rapid development of electronic commerce, Internet auctions have become increasingly popular over the years. [9, 15, 11]. Unlike traditional auctions, typical Internet auctions pose no geographical constraint. That is, sellers and bidders from all over the world can participate in an Internet auction remotely over the Internet, without having to physically attend the auction event. For sellers, this reduces the cost of running an auction. For bidders, this lowers the entry cost. Effectively, in an individually rational auction mechanism (a mechanism that guarantees nonnegative utilities for the agents), a bidder, at worst, loses nothing (but time) by participating in an auction. On the one hand, this encourages more bidders to join the auction, which potentially leads to higher revenue for the seller, as well as a higher social welfare for the bidders. On the other hand, it enables the bidders to manipulate by submitting multiple bids via multiple fictitious identities (*e.g.*, user accounts linked to different e-mail addresses).

The line of research on preventing manipulation via multiple fictitious identities in Internet auctions was explicitly framed by the groundbreaking work of Yokoo *et al.* [19]. 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 define an auction mechanism to be *false-name-proof* if the mechanism is not only strategy-proof, but also, under this mechanism, an agent cannot benefit from submitting multiple bids under false names (fictitious identities). The authors also extended the revelation principle [10] 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 [5]). These are the Set mechanism [16], the Minimal Bundle (MB) mechanism [16], and the Leveled Division Set (LDS) mechanism [18].[1] Other work on false-name-proofness includes the following. For general combinatorial auction settings, Yokoo [16] and Todo *et al.* [12] 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 [4, 14]. Finally, Conitzer [3] 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 beneficial to the agent's other identities, as shown in the following example:

Example 1. There are three single-minded[2] agents $1, 2, 3$ and

[1] A very recent paper [7] introduces a new mechanism called the ARP mechanism. However, this mechanism requires the additional restriction that agents are single-minded.

[2] A single-minded agent is only interested in a single, specific bun-

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 $\{A, B\}$ (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 benefit 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, resulting in a utility of 1.

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. Indeed, this paper also contributes to the research on false-name-proofness in the traditional sense. Since FNPW is stronger than FNP, the mechanisms we propose in this paper, as well as the automated mechanism design technique, should be of interest in the FNP context as well.

The paper is organized as follows. In Section 2, we formalize the problem we study. In Section 3, we give a sufficient and necessary condition on the type space for the VCG mechanism to be FNPW. In Section 4, we characterize both the payment rules and the allocation rules of FNPW mechanisms in general combinatorial auctions. We also derive a sufficient condition that can be used to check whether a mechanism is FNPW. In Section 5, we propose the *maximum marginal value item pricing (MMVIP)* mechanism, which we prove is FNPW. In Section 6, we propose an automated mechanism design technique that transforms any feasible mechanism into an FNPW mechanism. This technique builds on the sufficient condition in Section 4. In Section 7, we show that, under a minor condition, the mechanism that sells all the items as a single bundle has the highest worst-case efficiency ratio among all FNPW mechanisms. Finally, in the appendix, we give a characterization of FNP(W) social choice rules.

2. FORMALIZATION

We will use the following notation:

- $N = \{1, 2, \ldots, n\}$: the set of agents
- $G = \{1, 2, \ldots, m\}$: the set of items
- Θ: the type space of each agent
- $\theta_i \in \Theta$: agent i's reported type (since we consider only strategy-proof mechanisms, when there is no ambiguity, we also use θ_i to denote i's true type)
- $-i$: the set of agents other than agent i
- $\theta_{-i} \in \Theta^{n-1}$: types reported by agents other than agent i

We study combinatorial auction settings satisfying the following assumptions:

- Each agent has a *quasi-linear* utility function. That is, there exists a function v (determined by the setting) such that if an agent with true type $\theta \in \Theta$ ends up with bundle $B \subset G$ and payment $p \in \mathbb{R}$, then her utility equals $v(\theta, B) - p$.
- $\forall \theta \in \Theta$, we have $v(\theta, \emptyset) = 0$.
- $\forall B_1 \subseteq B_2 \subseteq G$, $\forall \theta \in \Theta$, we have $v(\theta, B_1) \leq v(\theta, B_2)$. That is, there is *free disposal*.
- An agent can have any valuation function satisfying the above conditions. That is, we are dealing with *rich domains* [1]. It should be noted that in Section 3, we study how restrictive the type space has to be in order for the VCG mechanism to be FNPW. That is, we do not have the rich-domain assumption in Section 3, which is an exception.

A mechanism consists of an allocation rule $X : (\Theta, \Theta^{n-1}) \to \mathcal{P}(G)$ and a payment rule $P : (\Theta, \Theta^{n-1}) \to \mathbb{R}$. $X(\theta_i, \theta_{-i})$ is the bundle agent i receives when reporting θ_i (when the other agents report θ_{-i}). $P(\theta_i, \theta_{-i})$ is the payment agent i has to make when reporting θ_i (when the other agents report θ_{-i}). When there is no ambiguity about the other agents' types, we simply use $X(\theta_i)$ and $P(\theta_i)$ in place of $X(\theta_i, \theta_{-i})$ and $P(\theta_i, \theta_{-i})$.

Throughout the paper, we only consider mechanisms satisfying the following conditions:

- *Strategy-proofness:* $\forall \theta_i, \theta_i', \theta_{-i}$, we have $v(\theta_i, X(\theta_i)) - P(\theta_i) \geq v(\theta_i, X(\theta_i')) - P(\theta_i')$. That is, if an agent uses only one identity, then truthful reporting is a dominant strategy.
- *Pay-only:* $\forall \theta_i, \theta_{-i}$, we have $P(\theta_i) \geq 0$.
- *Individual rationality:* $\forall \theta_i, \theta_{-i}$, we have $v(\theta_i, X(\theta_i)) - P(\theta_i) \geq 0$. That is, if an agent reports truthfully, then her utility is guaranteed to be nonnegative. This condition also implies that if an agent does not win any items, or has valuation 0 for all the items, then her payment must be 0.

dle of items.

- *Consumer sovereignty:* $\forall \theta_{-i}, \forall B \subseteq G$, there exists $\theta_i \in \Theta$ such that $X(\theta_i, \theta_{-i}) \supseteq B$. That is, no matter what the other agents bid, an agent can always win any bundle (possibly at the cost of a large payment).
- *Determinism and symmetry:* We only consider deterministic mechanisms that are symmetric over both the agents and the items (except for ties).

Yokoo [16] showed that in our setting, the mechanisms satisfying the above conditions coincide with the *(anonymous) price-oriented, rationing-free (PORF)* mechanisms. Similar price-based representations have also been presented by others, including [8]. The PORF mechanisms work as follows:

- The agents submit their reported types.
- The mechanism is characterized by a price function $\chi : \mathcal{P}(G) \times \Theta^{n-1} \rightarrow [0, \infty)$. For any agent i, for any multiset θ_{-i} of types reported by the other agents, for any set of items $S \subseteq G$, $\chi(S, \theta_{-i})$ is the price of S offered to i by the mechanism. That is, i can purchase S at a price of $\chi(S, \theta_{-i})$. $\forall \theta_{-i}$, we have $\chi(\emptyset, \theta_{-i}) = 0$. That is, the price of nothing is always zero. $\forall \theta_{-i}, \forall S_1 \subseteq S_2 \subseteq G$, we have $\chi(S_1, \theta_{-i}) \leq \chi(S_2, \theta_{-i})$. That is, a larger bundle always has a higher (or the same) price.
- The mechanism will select a bundle for agent i that is optimal for her given the prices, that is, the bundle chosen for i is in

  $$\arg\max_{S \subseteq G}\{v(\theta_i, S) - \chi(S, \theta_{-i})\}.$$

  The agent then pays the price for this bundle.
- Naturally, the mechanism must ensure that no item is allocated to two different agents. This involves setting prices carefully, as well as breaking ties.

Since all *feasible* mechanisms (mechanisms that satisfy the desirable conditions in our setting) are PORF mechanisms, besides using X (the allocation rule) and P (the payment rule) to refer to a mechanism, we can also use the price function χ to refer to a mechanism, namely, the PORF mechanism with price function χ.[3]

In the remainder of this section, we formally define the traditional false-name-proofness (FNP) condition, as well as our new false-name-proofness with withdrawal (FNPW) condition.

Definition 1. **FNP.** A mechanism characterized by allocation rule X and payment rule P is FNP if and only if it satisfies the following:

$$\forall \theta_i, \forall \theta_{i1}, \theta_{i2}, \ldots, \theta_{ik}, \forall \theta_{-i}, \text{ we have}$$
$$v(\theta_i, X(\theta_i, \theta_{-i})) - P(\theta_i, \theta_{-i}) \geq$$
$$v(\theta_i, \bigcup_{j=1}^{k} X(\theta_{ij}, \theta_{-i} \cup (\bigcup_{t \neq j} \theta_{it}))) - \sum_{j=1}^{k} P(\theta_{ij}, \theta_{-i} \cup (\bigcup_{t \neq j} \theta_{it}))$$

That is, truthful reporting using a single identifier is always better than submitting multiple false-name bids.

Definition 2. **FNPW.** A mechanism characterized by allocation rule X and payment rule P is FNPW if and only if it satisfies the following:

$$\forall \theta_i, \forall \theta_{i1}, \theta_{i2}, \ldots, \theta_{ik}, \forall \theta'_{i1}, \theta'_{i2}, \ldots, \theta'_{iq}, \forall \theta_{-i}, \text{ we have}$$
$$v(\theta_i, X(\theta_i, \theta_{-i})) - P(\theta_i, \theta_{-i}) \geq$$
$$v(\theta_i, \bigcup_{j=1}^{k} X(\theta_{ij}, \theta_{-i} \cup (\bigcup_{t \neq j} \theta_{it}) \cup (\bigcup \theta'_{it})))$$
$$- \sum_{j=1}^{k} P(\theta_{ij}, \theta_{-i} \cup (\bigcup_{t \neq j} \theta_{it}) \cup (\bigcup \theta'_{it}))$$

That is, truthful reporting using a single identifier is always better than submitting multiple false-name bids and then withdrawing some of them.

Actually, FNPW is exactly equivalent to FNP plus the following condition:

Definition 3. **Others' bids do not help (OBDNH).** A mechanism characterized by allocation rule X and payment rule P satisfies the OBDNH condition if and only if

$$\forall \theta_i, \forall \theta', \forall \theta_{-i}, \text{ we have}$$
$$v(\theta_i, X(\theta_i, \theta_{-i})) - P(\theta_i, \theta_{-i}) \geq$$
$$v(\theta_i, X(\theta_i, \theta_{-i} \cup \theta')) - P(\theta_i, \theta_{-i} \cup \theta')$$

That is, an agent's utility for reporting truthfully does not increase if we add another agent.

THEOREM 1. *FNPW is equivalent to FNP plus OBDNH.*

PROOF. We first prove that FNPW implies FNP and OBDNH.

It is straightforward that FNPW implies FNP. We only need to prove that FNPW implies OBDNH. $\forall \theta_i$, $\forall \theta_{-i}$, $\forall \theta'$, let $k = 1$, $\theta_{i1} = \theta_i$, $q = 1$, and $\theta'_{i1} = \theta'$. With these assignments, the FNPW condition reduces to the OBDNH condition.

We now prove that FNP and OBDNH together imply FNPW. $\forall \theta_i, \forall \theta'_{i1}, \theta'_{i2}, \ldots, \theta'_{iq}, \forall \theta_{-i}$, according to OBDNH, we have $v(\theta_i, X(\theta_i, \theta_{-i})) - P(\theta_i, \theta_{-i}) \geq v(\theta_i, X(\theta_i, \theta_{-i} \cup (\bigcup \theta'_{it}))) - P(\theta_i, \theta_{-i} \cup (\bigcup \theta'_{it}))$. Then, according to FNP, $\forall \theta_{i1}, \theta_{i2}, \ldots, \theta_{ik}$, replacing θ_{-i} by $\theta_{-i} \cup (\bigcup \theta'_{it})$, we obtain $v(\theta_i, X(\theta_i, \theta_{-i} \cup (\bigcup \theta'_{it}))) - P(\theta_i, \theta_{-i} \cup (\bigcup \theta'_{it})) \geq v(\theta_i, \bigcup_{j=1}^{k} X(\theta_{ij}, \theta_{-i} \cup (\bigcup_{t \neq j} \theta_{it}) \cup (\bigcup \theta'_{it}))) - \sum_{j=1}^{k} P(\theta_{ij}, \theta_{-i} \cup (\bigcup_{t \neq j} \theta_{it}) \cup (\bigcup \theta'_{it}))$. Combining the inequalities, we obtain exactly the FNPW condition. □

According to Theorem 1, to check whether an FNP mechanism is FNPW, we only need to check whether it satisfies OBDNH.

CLAIM 1. *The Leveled Division Set (LDS) mechanism [18] does not satisfy OBDNH. That is, LDS is not FNPW in general.*[4]

The general LDS mechanism is rather complicated. Instead of describing LDS in its general form, we focus on a specific LDS mechanism for three items, which is characterized by reserve price 1 and two levels: $[\{(A, B, C)\}]$ and $[\{(A, B), (C)\}, \{(A), (B, C)\}]$. The mechanism works as follows. If there are at least two agents whose valuations for $\{A, B, C\}$ are at least 3, then we combine $\{A, B, C\}$ into one bundle, and run the Vickrey auction. If every agent's valuation for $\{A, B, C\}$ is less than 3, then we do the following. We first introduce a dummy agent into the system. The dummy agent has an additive valuation function and values every item at 1. We only allow five types of allocations: 1) The dummy agent wins everything. 2) The dummy agent wins one of $\{A, B\}$

[3] Technically, there can be multiple PORF mechanisms with the same price function due to tie-breaking, but this will generally not be an issue.

[4] We will show later that the other two known FNP mechanisms, that is, the Set mechanism [16] and the Minimal Bundle mechanism [16], are both FNPW.

and $\{C\}$, and a non-dummy agent wins the other. 3) The dummy agent wins one of $\{A\}$ and $\{B, C\}$, and a non-dummy agent wins the other. 4) A non-dummy agent wins one of $\{A, B\}$ and $\{C\}$, and another non-dummy agent wins the other. 5) A non-dummy agent wins one of $\{A\}$ and $\{B, C\}$, and another non-dummy agent wins the other. We run the VCG mechanism on this restricted set of possible allocations. Finally, if there is only one agent whose valuation for $\{A, B, C\}$ is at least 3, then this agent is the only winner. She has the option to purchase all the items at price 3, or to obtain the result she would have obtained if everyone (including the dummy agent) were to join in the above maximal-in-range mechanism.

PROOF. We only need to prove that the above specific LDS mechanism does not satisfy OBDNH. We consider the following scenario. There are two agents. Agent 1 bids 2.2 on $\{A, B\}$. Agent 2 is single-minded, valuing $\{A\}$ at 1.1. Under the above LDS mechanism, if 2 reports truthfully, then $\{A, B\}$ is allocated to 1, and $\{C\}$ is allocated to the dummy agent (thrown away). That is, if 2 reports truthfully, then her utility equals 0. If, besides 2's true identity, 2 also submits a false-name bid of 2.9 on $\{B, C\}$, then $\{B, C\}$ will be allocated to 2's false-name identity (2 will withdraw this identity, that is, refuse to pay for this bundle), and $\{A\}$ will be allocated to 2's true identity at a price of 1. That is, 2 now has utility 0.1. We conclude that, in general, LDS does not satisfy OBDNH, and hence is not FNPW. □

3. RESTRICTION ON THE TYPE SPACE SO THAT VCG IS FNPW

The VCG mechanism [13, 2, 6] satisfies several nice properties, including efficiency, strategy-proofness, individual rationality, and the non-deficit property. Unfortunately, as shown by Yokoo *et al.* [19], the VCG mechanism is not FNP for general type spaces. One sufficient condition on the type space for the VCG mechanism to be FNP is as follows:

Definition 4. **Submodularity [19].** For any set of bidders Y, whose types are drawn from Θ, $\forall S_1, S_2 \subseteq G$, 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 defined as the total utility of bidders in Y, if we allocate items in S to these bidders efficiently.

That is, if the type space Θ satisfies the above condition, then the VCG mechanism is FNP. In this section, we aim to characterize type spaces for which VCG is FNPW. We consider restricted type spaces (that make the VCG mechanism FNPW) in this section. In other sections, unless specified, we assume that the rich-domain condition holds.

THEOREM 2. *If the type space satisfies the submodularity condition, then the VCG mechanism is FNPW. Conversely, if the mechanism is FNPW, and additionally the type space contains the additive valuations, then the type space satisfies the submodularity condition.*

That is, submodularity does not only imply FNP, it actually implies FNPW. Moreover, unlike for FNP, in the case of FNPW, the converse also holds—if we allow the additive valuations (those valuations which value any set of items at the sum of the values of its elements, with no complementarity and no substitutability).

PROOF. We first prove that if the type space satisfies submodularity, then the VCG mechanism is FNPW. We consider agent i. Let K be the set of false-name identities i submits and keeps at the end. Let W be the set of false-name identities i submits and withdraws. We already know that submodularity is sufficient for the VCG mechanism to be FNP. Hence, if K contains multiple identities, then i might as well replace all of them by one identity that reports i's true type. We then show that the identities in W do not help i (OBDNH). We use S to denote the set of items won by i at the end. To win S, i pays the VCG price $U(G, \{-i\} \cup W) - U(G - S, \{-i\} \cup W)$ ($\{-i\}$ is the set of other agents). We use S' to denote the set of items won by identities in W, when we allocate items in $G - S$ to identities in $\{-i\} \cup W$ efficiently. We have that $U(G, \{-i\} \cup W) - U(G - S, \{-i\} \cup W) = U(G, \{-i\} \cup W) - U(G - S - S', \{-i\}) - U(S', W) \geq U(G - S', \{-i\}) + U(S', W) - U(G - S - S', \{-i\}) - U(S', W) = U(G - S', \{-i\}) - U(G - S - S', \{-i\})$. The submodularity condition implies that $U(G - S', \{-i\}) - U(G - S - S', \{-i\}) \geq U(G, \{-i\}) - U(G - S, \{-i\})$. But, the expression on the right-hand side of the inequality is the price i would be charged for S when she uses a single identifier. That is, i does not benefit from the false-name identities in W. Therefore, the VCG mechanism is FNPW if the type space satisfies submodularity.

Next, we prove that if the VCG mechanism is FNPW, then the type space must satisfy submodularity (if it contains the additive valuations). Let S be an arbitrary set of items. Let i be an agent that is interested in S. Since we allow additive valuations, such i always exists (*e.g.*, i may have a very large valuation for every item in S). If i bids truthfully, then she can win S at a price of $U(G, \{-i\}) - U(G - S, \{-i\})$. Let S' be another arbitrary set of items that does not intersect with S. For each item j in S', we introduce a false-name identity that is only interested in item j, with value c, where c is set to a very large value (*e.g.*, larger than $U(G, \{-i\})$). These false-name identities are allowed since we assume the type space contains the additive valuations. Let W be the set of identities introduced. With W, i can win S at a price of $U(G, \{-i\} \cup W) - U(G - S, \{-i\} \cup W)$. We have that $U(G, \{-i\} \cup W) - U(G - S, \{-i\} \cup W) = U(G - S', \{-i\}) + U(S', W) - U(G - S - S', \{-i\}) - U(S', W) = U(G - S', \{-i\}) - U(G - S - S', \{-i\})$. The new price should never be smaller than the old price. Otherwise, there is an incentive for i to submit false-name bids and withdraw them. That is, we have $U(G, \{-i\}) - U(G - S, \{-i\}) \leq U(G - S', \{-i\}) - U(G - S - S', \{-i\})$. Let $S_1 = G - S$, $S_2 = G - S'$, and $Y = \{-i\}$. We have $U(S_1 \cap S_2, Y) - U(S_1, Y) \leq U(S_2, Y) - U(S_1 \cup S_2, Y)$. Since S_1, S_2, and Y are arbitrary, we have submodularity. □

4. CHARACTERIZATION OF FNPW MECHANISMS

Yokoo [16] and Todo *et al.* [12] characterized the payment rules (the price functions in the PORF representation) and the allocation rules of FNP mechanisms, respectively. In this section, we present similar results on the characterization of FNPW mechanisms.

4.1 Characterizing FNPW payments

We recall that in our setting, a feasible mechanism corresponds to a PORF mechanism, characterized by a price function χ. Yokoo [16] gave the following sufficient and necessary condition on χ for the mechanism characterized by χ to be FNP.

Definition 5. **No superadditive price increase (NSA).** Let O be an arbitrary set of agents.[5] We run mechanism χ (a PORF mechanism characterized by price function χ) for the agents in O. Let

[5]In a slight abuse of language, we also use "a set of agents" to refer to the types reported by this set of agents.

Y be an arbitrary subset of O. Let B_i ($i \in Y$) be the set of items agent i obtains. We must have

$$\sum_{i \in Y} \chi(B_i, O - \{i\}) \geq \chi(\bigcup_{i \in Y} B_i, O - Y).$$

By modifying the NSA condition, we get the following sufficient and necessary condition on χ for mechanism χ to be FNPW.

Definition 6. **No superadditive price increase with withdrawal (NSAW).** Let O be an arbitrary set of agents. We run mechanism χ for the agents in O. Let Y and Z be two arbitrary nonintersecting subsets of O. Let B_i ($i \in Y$) be the set of items agent i obtains. We must have

$$\sum_{i \in Y} \chi(B_i, O - \{i\}) \geq \chi(\bigcup_{i \in Y} B_i, O - Y - Z).$$

NSAW is equivalent to NSA plus the following condition.

Definition 7. **Prices increase with agents (PIA).** Let O be an arbitrary set of agents. Let a be another agent not in O. $\forall S \subseteq G$, we must have

$$\chi(S, O \cup \{a\}) \geq \chi(S, O).$$

That is, from the perspective of agent i, if another agent joins in, then the price i faces for any set of items must (weakly) increase.

CLAIM 2. *NSAW is equivalent to NSA plus PIA.*

PROOF. We first prove that NSAW implies NSA and PIA. It is straightforward that NSAW implies NSA, so we only need to show that NSAW implies PIA. Let R be an arbitrary set of agents. Let a be another agent not in R. $\forall S \subseteq G$, we can construct an agent (denoted by y) that wins S if we run χ on the agents in $R \cup \{a\} \cup \{y\}$ (*e.g.*, let y be single-minded on S, with a very large value). Let $Y = \{y\}$, $Z = \{a\}$, and $O = R \cup Y \cup Z$. NSAW implies that $\chi(S, R \cup Z) = \chi(S, R \cup \{a\}) \geq \chi(S, R)$. That is, NSAW implies PIA.

We now prove that NSA and PIA imply NSAW. PIA implies that $\chi(\bigcup_{i \in Y} B_i, O - Y - Z) \leq \chi(\bigcup_{i \in Y} B_i, O - Y)$. NSA implies that $\sum_{i \in Y} \chi(B_i, O - \{i\}) \geq \chi(\bigcup_{i \in Y} B_i, O - Y)$. Combining the two inequalities, we obtain the NSAW condition. □

THEOREM 3. *Mechanism χ is FNPW if and only if χ satisfies the NSAW condition.*

PROOF. We first prove that if χ satisfies NSAW, then the mechanism is FNPW. Let us consider a specific agent x. Let $O - Y - Z$ be the set of agents other than herself. Let Y be the set of false-name identities x submits and keeps at the end. Let Z be the set of false-name identities x submits but withdraws at the end. So, O is the set of all the identities. The set of items x receives at the end is $\bigcup_{i \in Y} B_i$, where B_i is the bundle won by identity i. The total price x pays is $\sum_{i \in Y} \chi(B_i, O - \{i\})$. According to NSAW, this price is at least $\chi(\bigcup_{i \in Y} B_i, O - Y - Z)$. That is, x would not be any worse off if she just used a single identity to buy $\bigcup_{i \in Y} B_i$. When x uses only one identity, her optimal strategy is to report truthfully. Therefore, if NSAW is satisfied, mechanism χ is FNPW.

Next, we prove that if mechanism χ is FNPW, then χ must satisfy NSAW. Suppose not, that is, suppose there exists some χ that corresponds to an FNPW mechanism, and there exist three nonintersecting sets of agents Y, Z, and $O - Y - Z$, such that $\sum_{i \in Y} \chi(B_i, O - \{i\}) < \chi(\bigcup_{i \in Y} B_i, O - Y - Z)$, where B_i is the bundle agent i obtains (when we apply mechanism χ to the agents in O). Let us consider a single-minded agent x, who values $\bigcup_{i \in Y} B_i$ at exactly $\chi(\bigcup_{i \in Y} B_i, O - Y - Z)$. If the set of other agents faced by x is $O - Y - Z$, then x has utility 0 if she reports truthfully using a single identifier. However, if x instead submits multiple false-name identities $Y + Z$, keeps those in Y and withdraws those in Z, then she will obtain her desired items at a lower price and end up with positive utility, contradicting the assumption that χ is FNPW. That is, if NSAW is not satisfied, then χ is not FNPW. □

4.2 A sufficient condition for FNPW

The NSAW condition in Section 4.1 leads to the following sufficient condition for mechanism χ to be FNPW.

Definition 8. **Sufficient condition for no superadditive price increase with withdrawal (S-NSAW).** Let O be an arbitrary set of agents. S-NSAW holds if we have both of the following conditions:

- **Discounts for larger bundles (DLB).** $\forall S_1, S_2 \subseteq G$ with $S_1 \cap S_2 = \emptyset$, $\chi(S_1, O) + \chi(S_2, O) \geq \chi(S_1 \cup S_2, O)$. That is, the sum of the prices of two disjoint sets of items must be at least the price of the joint set.
- **Prices increase with agents (PIA).**[6] $\forall S \subseteq G$, for any agent a that is not in O, $\chi(S, O \cup \{a\}) \geq \chi(S, O)$.

CLAIM 3. *Mechanism χ is FNPW if χ satisfies S-NSAW.*

PROOF. We only need to show that S-NSAW is stronger than NSAW (by Theorem 3, NSAW is sufficient (and necessary) for χ to be FNPW). Let χ satisfy S-NSAW. Let O be an arbitrary set of agents. We run mechanism χ on the agents in O. We divide O into three subgroups, Y, Z, and $O - Y - Z$. For $i \in Y$, let B_i be the bundle agent i obtains. By PIA, we have $\sum_{i \in Y} \chi(B_i, O - \{i\}) \geq \sum_{i \in Y} \chi(B_i, O - Y - Z)$. By DLB, we have $\sum_{i \in Y} \chi(B_i, O - Y - Z) \geq \chi(\bigcup_{i \in Y} B_i, O - Y - Z)$. Combining these inequalities, we can conclude that S-NSAW implies NSAW. □

S-NSAW is a cleaner, but more restrictive condition than NSAW. (To see why, note that even if DLB does not hold, NSA may still hold: even if $\chi(S_1, O) + \chi(S_2, O) < \chi(S_1 \cup S_2, O)$, it may be the case that by putting separate bids on S_1 and S_2, each of these bids makes the price for the other bundle go up, so that the result is still more expensive than buying $S_1 \cup S_2$ as a single bundle.) We find it easier to use S-NSAW to prove that a mechanism is FNPW (rather than using the more complex NSAW condition).[7] Let us recall the three existing FNP mechanisms (for general combinatorial auction settings): the Set mechanism, the MB Mechanism, and the LDS mechanism. We have already shown that LDS is not FNPW. With the help of S-NSAW, we can prove that both Set and MB are FNPW.

CLAIM 4. *Both the Set mechanism and the MB mechanism satisfy the S-NSAW condition. Hence, they are FNPW.*

The Set mechanism simply combines all the items into a grand bundle. The grand bundle is then sold in a Vickrey auction. The MB (Minimal Bundle) mechanism builds on the concept of minimal bundles. A set of items S ($\emptyset \subsetneq S \subset G$) is called a *minimal*

[6] This is the same PIA condition as the one in Section 4.1.

[7] However, S-NSAW cannot be used to prove that a mechanism is *not* FNPW, because it is a more restrictive condition.

bundle for agent i if and only if $\forall S' \subsetneq S$, $v(i,S) > v(i,S')$. Under the MB mechanism, the price of a bundle S an agent faces is equal to the highest valuation value of a bundle, which is minimal and conflicting with S. Generally, MB coincides with Set, because usually the grand bundle is a minimal bundle for every agent (any smaller bundle usually gives at least slightly lower utility).

PROOF. The proof of the above claim is straightforward; we omit the details due to space constraint. □

We will also use S-NSAW to prove that the MMVIP mechanism that we propose (Section 5) is FNPW. The automated mechanism design technique for generating FNPW mechanisms (Section 6) is also based on S-NSAW.

4.3 Characterizing FNPW allocations

Todo *et al.* [12] gave the following characterization of the allocation rules of FNP mechanisms. We recall that $X(\theta_i, \theta_{-i})$ is the set of items that agent i wins if her reported type is θ_i and the reported types of the other agents are θ_{-i}. To simplify notation, we use $X(\theta_i)$ in place of $X(\theta_i, \theta_{-i})$ when there is no risk of ambiguity.

Definition 9. **Weak-monotonicity [1].** X is weakly monotone if $\forall \theta_i, \theta_i', \theta_{-i}$, we have

$$v(\theta_i, X(\theta_i)) - v(\theta_i, X(\theta_i')) \geq v(\theta_i', X(\theta_i)) - v(\theta_i', X(\theta_i')).$$

Definition 10. **Sub-additivity [12].** $\forall \theta_i, \forall \theta_i', \forall \theta_{i1}, \theta_{i2}, \ldots, \theta_{ik}$, $\forall \theta_{i1}', \theta_{i2}', \ldots, \theta_{ik}'$, $\forall \theta_{-i}$, we have the following:

$$\begin{gathered} X(\theta_i) = \bigcup_{l=1}^{k} X_{+I_{-l}^k}(\theta_{il}) \\ v(\theta_i', X(\theta_i')) = 0 \\ X_{+I_{-l}^k}(\theta_{il}') \supseteq X_{+I_{-l}^k}(\theta_{il}) \\ v(\theta_{il}', X_{+I_{-l}^k}(\theta_{il}')) = v(\theta_{il}', X_{+I_{-l}^k}(\theta_{il})) \\ \Downarrow \\ v(\theta_i', X(\theta_i)) \leq \sum_{l=1}^{k} v(\theta_{il}', X_{+I_{-l}^k}(\theta_{il})). \end{gathered}$$

(Here, $X_{+I_{-l}^k}(\theta_{il})$ is short for $X(\theta_{il}, \theta_{-i} \cup (\bigcup_{1 \leq t \leq k, t \neq l} \theta_{it}))$.)

X is said to be *FNP-implementable* if there exists a payment rule P so that X combined with P constitutes a feasible FNP mechanism. Todo *et al.* [12] showed that X is FNP-implementable if and only X satisfies both weak-monotonicity and sub-additivity.

We define allocation rule X to be *FNPW-implementable* if there exists a payment rule P so that X combined with P constitutes a feasible FNPW mechanism. We introduce a third condition called *withdrawal-monotonicity*. We prove that X is FNPW-implementable if and only X satisfies weak-monotonicity, sub-additivity, and withdrawal-monotonicity.

Definition 11. **Withdrawal-monotonicity.** $\forall \theta_i, \forall \theta_{-i}, \forall \theta^a, \forall \theta_i^L$, $\forall \theta_i^U$, the following holds:

$$\begin{gathered} v(\theta_i^L, X(\theta_i^L, \theta_{-i})) = 0 \\ X(\theta_i^U, \theta_{-i} \cup \theta^a) = X(\theta_i, \theta_{-i}) \\ \Downarrow \\ v(\theta_i^L, X(\theta_i, \theta_{-i})) \leq v(\theta_i^U, X(\theta_i, \theta_{-i})) \end{gathered}$$

THEOREM 4. *An allocation rule X is FNPW-implementable if and only X satisfies weak-monotonicity, sub-additivity, and withdrawal-monotonicity.*

PROOF. We first prove that if X is FNPW-implementable, then X satisfies weak-monotonicity, sub-additivity, and withdrawal-monotonicity. If X is FNPW-implementable, then X is also FNP-implementable. Hence, X satisfies both weak-monotonicity and sub-additivity [12]; only withdrawal-monotonicity remains to be shown. Let χ be the (PORF) price function corresponding to an FNPW mechanism that allocates according to X. We denote $X(\theta_i, \theta_{-i})$ by S. Since $v(\theta_i^L, X(\theta_i^L, \theta_{-i})) = 0$, we have $v(\theta_i^L, S) \leq \chi(S, \theta_{-i})$ (otherwise, an agent with true type θ_i^L would be better off purchasing S). Since $X(\theta_i^U, \theta_{-i} \cup \theta^a) = X(\theta_i, \theta_{-i}) = S$, we have $v(\theta_i^U, S) \geq \chi(S, \theta_{-i} \cup \theta^a)$ (because an agent with true type θ_i^U is best off buying S when the other agents' types are $\theta_{-i} \cup \theta^a$). χ is FNPW, hence it satisfies the PIA condition, by Theorem 3 and Claim 2. So, we have $\chi(S, \theta_{-i} \cup \theta^a) \geq \chi(S, \theta_{-i})$. Combining all the inequalities, we get $v(\theta_i^U, X(\theta_i, \theta_{-i})) \geq v(\theta_i^L, X(\theta_i, \theta_{-i}))$. That is, withdrawal-monotonicity is satisfied.

Next, we prove that if X satisfies weak-monotonicity, sub-additivity, and withdrawal-monotonicity, then X is FNPW-implementable. Since X satisfies both weak-monotonicity and sub-additivity, X is FNP-implementable [12]. Let χ be a (PORF) price function that characterizes an FNP mechanism that allocates according to X. We prove that χ must also be FNPW. We only need to prove that χ satisfies PIA (because, according to Claim 2 and Theorem 3, if an FNP mechanism satisfies PIA, then it is FNPW). Suppose χ does not satisfy PIA. Then, there exists a set of agents O, an agent a not in O (where a's type is denoted by θ^a), and some $S \subseteq G$, such that $\chi(S, O) > \chi(S, O \cup \{a\})$. Let $\chi(S, O) - \chi(S, O \cup \{a\}) = \beta > 0$. Let θ_{-i} be the reported types of the agents in O. Let i be an agent that is single-minded on S, with a very large valuation, so that $X(\theta_i, \theta_{-i}) = S$ (we denote agent i's type by θ_i). We also construct an agent that is single-minded on S, with valuation $\chi(S, O) - \frac{\beta}{3}$. We denote the type of this agent by θ_i^L. We have $X(\theta_i^L, \theta_{-i}) = \emptyset$ (she is not willing to pay $\chi(S, O)$ to purchase S). Hence, $v(\theta_i^L, X(\theta_i^L, \theta_{-i})) = 0$. We construct another agent that is also single-minded on S, with valuation $\chi(S, O \cup \{a\}) + \frac{\beta}{3}$. We denote the type of this agent by θ_i^U. We have $X(\theta_i^U, \theta_{-i} \cup \theta^a) = S = X(\theta_i, \theta_{-i})$. By withdrawal-monotonicity, we must have $v(\theta_i^L, X(\theta_i, \theta_{-i})) \leq v(\theta_i^U, X(\theta_i, \theta_{-i}))$. However, on the other hand, $v(\theta_i^L, X(\theta_i, \theta_{-i})) = \chi(S, O) - \frac{\beta}{3} = \chi(S, O \cup \{a\}) + \frac{2\beta}{3} > \chi(S, O \cup \{a\}) + \frac{\beta}{3} = v(\theta_i^U, X(\theta_i, \theta_{-i}))$. We have reached a contradiction. We conclude that χ has to satisfy PIA, which implies that χ is FNPW. Hence, X is FNPW-implementable. □

5. MAXIMUM MARGINAL VALUE ITEM PRICING MECHANISM

In this section, we introduce a new FNPW mechanism.

Definition 12. **Maximum marginal value item pricing mechanism (MMVIP).** Let O be an arbitrary set of agents. MMVIP is characterized by the following price function χ.

- $\forall S \subseteq G$, $\chi(S, O) = \sum_{s \in S} \chi(\{s\}, O)$. That is, χ uses *item pricing*.
- $\forall s \in G$, $\chi(s, O) = \max_{j \in O} \max_{S \subseteq G - \{s\}} \{v(j, S \cup \{s\}) - v(j, S)\}$.[8] That is, 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.

[8]In this notation, we assume that the maximum over an empty set is 0 (for presentation purpose). Such notation will also appear later in the paper.

CLAIM 5. *MMVIP is feasible and FNPW.*

PROOF. We first prove that MMVIP is feasible. We need to show that, with appropriate tie-breaking, MMVIP will never allocate the same item to multiple agents. Let us suppose that under MMVIP there is a scenario in which two agents, i and j, both win item s. Let S_i and S_j be the sets of other items (items other than s) that i and j win at the end, respectively. Let $v_i = v(i, S_i \cup \{s\}) - v(i, S_i)$. That is, v_i is i's marginal value for s. Let $v_j = v(j, S_j \cup \{s\}) - v(j, S_j)$. That is, v_j is j's marginal value for s. If $v_i > v_j$, then j has to pay at least v_i to win s, which is too high for her; j is better off not winning s. Similarly, if $v_i < v_j$, then i is better off not winning s. If $v_i = v_j$, then i and j both have to pay at least their marginal value for s to win s. That is, they are either indifferent between winning s or not, or prefer not to win. The only case that does not lead to a contradiction is where they are both indifferent; any tie-breaking rule can resolve this conflict.

We then show that MMVIP is FNPW. By Claim 3, we only need to prove that the price function χ that characterizes MMVIP satisfies S-NSAW. Let O be an arbitrary set of agents. $\forall S_1, S_2 \subseteq G$ with $S_1 \cap S_2 = \emptyset$, we have $\chi(S_1, O) + \chi(S_2, O) = \chi(S_1 \cup S_2, O)$, because MMVIP uses item pricing. Hence, DLB is satisfied. $\forall S \subseteq G$, for any agent a that is not in O, $\chi(S, O \cup \{a\}) = \sum_{s \in S} \chi(s, O \cup \{a\}) = \sum_{s \in S} \max_{j \in O \cup \{a\}} \max_{S' \subseteq G - \{s\}} \{v(j, S' \cup \{s\}) - v(j, S')\} \geq \sum_{s \in S} \max_{j \in O} \max_{S' \subseteq G - \{s\}} \{v(j, S' \cup \{s\}) - v(j, S')\} = \sum_{s \in S} \chi(s, O) = \chi(S, O)$. That is, PIA is also satisfied. □

Next, we prove two properties of the MMVIP mechanism.

CLAIM 6. *Suppose we restrict the domain to additive valuations. Then, MMVIP coincides with the VCG mechanism, so that MMVIP=VCG is FNPW and efficient. Moreover, no efficient, strategy-proof, and individually rational mechanism achieves strictly higher revenue than MMVIP=VCG on any additive profile.*

PROOF. When the agents' valuations are additive, we have that MMVIP's item price function satisfies $\chi(s, O) = \max_{j \in O} \max_{S \subseteq G - \{s\}} \{v(j, S \cup \{s\}) - v(j, S)\} = \max_{j \in O} v(j, \{s\})$. Thus, MMVIP is equivalent to m separate Vickrey auctions (one Vickrey auction for each item), and hence to VCG (which also corresponds to m separate Vickrey auctions when the valuations are additive).

Now, for the sake of contradiction, let us assume that there exists an efficient, strategy-proof, and individually rational mechanism M that achieves higher revenue than MMVIP=VCG on at least one additive profile. Let χ' be the price function of M. Let χ be the price function of MMVIP. If $\chi'(S, O) \leq \chi(S, O)$ for every set of items S and every set of additive agents O, then MMVIP's revenue is at least that of M on every additive profile, because due to the efficiency of both mechanisms, we can assume without loss of generality that the winners and the bundles won by the winners are the same under both mechanisms. Hence, it must be the case that there exists a set of items S and a set of additive agents O, so that $\chi'(S, O) > \chi(S, O)$. Let $x_s = \max_{j \in O} v(j, \{s\})$. If an additive agent's valuation for s is $x_s + \epsilon$ for all $s \in S$, then her valuation for S is $\chi(S, O) + |S|\epsilon$. Hence, for sufficiently small ϵ, $\chi(S, O) + |S|\epsilon < \chi'(S, O)$, so this agent will not win S against O under M. But this contradicts the assumed efficiency of M. Hence, the claim holds. □

The above claim essentially says that, when the agents' valuations are additive, MMVIP "does the right thing." MMVIP is the only known FNP/FNPW mechanism with the above property for general combinatorial auctions.

Before moving on to the other property that we prove about MMVIP, we first experimentally compare the revenue and allocative efficiency of the MMVIP mechanism and the Set mechanism, under the assumption that the agents' valuations are additive.[9] We assume that there are 100 items and 100 agents. An agent's valuation for an item is drawn i.i.d. from $U(0, 1)$ (the uniform distribution from 0 to 1). The results are presented below (the numbers shown are average over 10000 instances):

	MMVIP	Set
Revenue	98.02	56.19
Efficiency	99.01	57.22

Finally, we have the following claim about MMVIP.

CLAIM 7. *Among all FNPW mechanisms that use item pricing, MMVIP has minimal payments. That is, let χ be the price function of MMVIP. Let χ' be a different price function corresponding to a different FNPW mechanism M that also uses item pricing. We have that there always exists a set of items S and a set of agents O, so that $\chi'(S, O) > \chi(S, O)$.*

PROOF. For the sake of contradiction, let us assume that the claim is false. That is, we assume that for every set of items S and every set of agents O, we have $\chi'(S, O) \leq \chi(S, O)$. Since $\chi \neq \chi'$, we have that there exists at least one set of items S and one set of agents O such that $\chi'(S, O) < \chi(S, O)$. Since $\chi'(S, O) = \sum_{s \in S} \chi'(s, O)$ and $\chi(S, O) = \sum_{s \in S} \chi(s, O)$, it follows that there exists $s \in S$ such that $\chi'(s, O) < \chi(s, O)$. By the definition of MMVIP, $\chi(s, O)$ corresponds to the maximal marginal value of some agent $j \in O$. That is, there exists $S' \subset G$ with $s \notin S'$ such that $\chi(s, O) = v(j, S' \cup \{s\}) - v(j, S')$. We construct an agent x, whose valuation function is additive. Let x's valuations of items not in $S' \cup \{s\}$ be extremely high, so that x wins all these items under both mechanisms χ and χ'. (We recall that we assume consumer sovereignty for FNPW mechanisms, so that $\chi, \chi' < \infty$ everywhere.) Let x's valuation on s be $\chi(s, O) - \epsilon$ (where ϵ is small enough so that $\chi(s, O) - \epsilon > \chi'(s, O)$). Let x's valuation of items in S' be 0. When the set of agents consists of x and the agents in O, we have that x wins all the items except for those in S' under M. Since M is FNPW, we have $\chi'(s, O) \geq \chi'(s, \{j\})$ (PIA). That is, when the set of agents consists of only x and j, x also wins all the items except for those in S' under M. Also, under M, j wins all of S', because for any $s' \in S'$, we have $\chi'(s', \{x\}) \leq \chi(s', \{x\}) = 0$. However, we then have that $\chi'(s, \{x\}) \leq \chi(s, \{x\}) = \chi(s, O) - \epsilon = v(j, S' \cup \{s\}) - v(j, S') - \epsilon$, so that j would choose to also win s when facing x under M. That is, under M, when the set of agents consists of only x and j, s is won by both agents, contradicting the assumption that M is feasible. Thus, assuming that the claim is false leads to a contradiction. □

6. AUTOMATED FNPW MECHANISM DESIGN

In this section, we propose an automated mechanism design (AMD) technique that transforms any feasible mechanism into an FNPW mechanism. In our setting, a feasible mechanism is characterized by a price function χ. We start with any χ that corresponds

[9]Under this assumption, the VCG mechanism coincides with the MMVIP mechanism. We also have that the MB mechanism and the Set mechanism coincide. (In our experimental setup, the grand bundle is always a minimal bundle for every agent.)

to a feasible mechanism (*e.g.*, the price function of the VCG mechanism). Our technique modifies χ so that it satisfies S-NSAW, while maintaining feasibility.

We recall that for general combinatorial auction settings, there are three known FNPW mechanisms (Set, MB, and MMVIP), and four known FNP mechanisms (the aforementioned three mechanisms, plus LDS). Though computationally expensive (like many other AMD techniques in other contexts), this technique has the potential to enlarge the set of known FNPW (FNP) mechanisms. By designing tiny instances of FNPW mechanisms via automated mechanism design, we may get a better understanding of the structure of FNPW mechanisms, from which we can then conjecture FNPW mechanisms in analytical form. Later in this section, we show that in a specific setting, by starting with the VCG mechanism, the AMD technique produces exactly the MMVIP mechanism. That is, had we not known the MMVIP mechanism, the AMD technique could have helped us find it (though it just so happened that we discovered MMVIP before the AMD technique). It remains an open question of whether new, general FNPW mechanisms can be found in this way.

6.1 Technique description

Let $H : \Theta^k \to [0, \infty)$ be a function that maps any set of agents O (more precisely, their reported types) to a nonnegative number $H(O)$. For any feasible mechanism χ, we define χ^H as follows:

- For any set of agents O, $\forall \emptyset \subsetneq S \subseteq G$, $\chi^H(S, O) = \chi(S, O) + H(O)$.
- For any set of agents O, $\chi^H(\emptyset, O) = \chi(\emptyset, O) = 0$.

That is, moving from χ to χ^H, if we fix the reported types of the other agents O, then we are essentially increasing the price of every nonempty set of items by the same amount, while keeping the price of $\emptyset$ at 0.

LEMMA 1. *[17] $\forall$ feasible χ, $\forall H$, χ^H is feasible.*

This lemma was first proved in [17].[10] An agent is allocated her favorite set of items (the set that maximizes valuation minus payment) in (PORF) mechanism χ. From the perspective of agent i, the set of types reported by the other agents θ_{-i} is fixed. That is, for i, under χ^H, the price of every nonempty set of items is increased by the same amount $H(\theta_{-i})$. Hence, agent i's favorite set of items is either unchanged, or has become $\emptyset$ (if $H(\theta_{-i})$ is too large). It is thus easy to see that if χ never allocates the same item to more than one agent, then neither does χ^H. That is, feasibility is not affected.[11]

THEOREM 5. *$\forall$ feasible χ, we define the following H. For any set of agents O, $H(O)$ equals the maximum of the following two values:*

- $\max_{S_1, S_2 \subseteq G, S_1 \cap S_2 = \emptyset} \{\chi(S_1 \cup S_2, O) - \chi(S_1, O) - \chi(S_2, O)\}$
- $\max_{\emptyset \subsetneq S \subseteq G, j \in O} \{\chi(S, O - \{j\}) + H(O - \{j\}) - \chi(S, O)\}$

We have that χ^H is FNPW.

It should be noted that, for any O, the first expression in the theorem is at least 0 (setting $S_1 = S_2 = \emptyset$). That is, H never takes negative values. χ^H is feasible by Lemma 1.

PROOF. We prove that χ^H satisfies S-NSAW. By Claim 3, this suffices to show that χ^H is FNPW.

Proof of DLB: Let O be an arbitrary set of agents. $\forall S_1, S_2 \subseteq G$ with $S_1 \cap S_2 = \emptyset$, we prove that $\chi^H(S_1, O) + \chi^H(S_2, O) \geq \chi^H(S_1 \cup S_2, O)$. If at least one of S_1 and S_2 is empty, then w.l.o.g., we assume $S_1 = \emptyset$. In this case, $\chi^H(S_1, O) + \chi^H(S_2, O) = \chi^H(S_2, O) = \chi^H(S_1 \cup S_2, O)$. If neither S_1 nor S_2 is empty, then we have $\chi^H(S_1, O) + \chi^H(S_2, O) - \chi^H(S_1 \cup S_2, O) = H(O) + \chi(S_1, O) + \chi(S_2, O) - \chi(S_1 \cup S_2, O) \geq H(O) - \max_{S_1' \cap S_2' = \emptyset} \{\chi(S_1' \cup S_2', O) - \chi(S_1', O) - \chi(S_2', O)\} \geq 0$.

Proof of PIA: Let O be an arbitrary set of agents. Let a be an agent that is not in O. If S is empty, then we have $\chi^H(S, O \cup \{a\}) = \chi^H(S, O) = 0$. $\forall \emptyset \subsetneq S \subseteq G$, $\chi^H(S, O \cup \{a\}) = H(O \cup \{a\}) + \chi(S, O \cup \{a\}) \geq (\chi(S, O) + H(S, O) - \chi(S, O \cup \{a\})) + \chi(S, O \cup \{a\}) = \chi^H(S, O)$. $\square$

This still leaves the question of how to compute the H described in the theorem; we address this next. Given χ, for any agent i and any set of other types θ_{-i}, we compute $H(\theta_{-i})$ using the following dynamic programming algorithm.

```
For t = 0, 1, ..., |θ_{-i}|
    For any T ⊆ θ_{-i} with |T| = t
        h_1 = max_{S_1, S_2 ⊆ G, S_1 ∩ S_2 = ∅} {χ(S_1 ∪ S_2, T) − χ(S_1, T) − χ(S_2, T)}.
        h_2 = max_{∅ ⊊ S ⊆ G, j ∈ T} {H(T − {j}) + χ(S, T − {j}) − χ(S, T)}.
        H(T) = max{h_1, h_2}.
```

CLAIM 8. *If we apply the AMD technique to a mechanism that already satisfies S-NSAW, the mechanism remains unchanged.*

We use the phrase "the AMD mechanism" to denote the mechanism generated by the AMD technique starting from VCG (though the AMD technique is not restricted to starting from VCG). Next, we prove a claim that is similar to Claim 6.

CLAIM 9. *When we restrict the preference domain to additive valuations, the MMVIP, the VCG, and the AMD mechanism all coincide.*

PROOF. Claim 6 already shows that MMVIP and VCG coincide. All that remains to show is that VCG already satisfies S-NSAW, so that by Claim 8, AMD is also the same. When the agents' valuations are additive, the VCG mechanism's price function χ is defined as follows: for any set of items $S \subset G$ and any set of additive agents O, $\chi(S, O) = \sum_{s \in S} x^s$, where x^s is the highest valuation for item s among the agents in O. It is easy to see that χ satisfies S-NSAW. $\square$

[10] The GM-SMA mechanism [17] relies on this property. However, it has been shown that GM-SMA is *not* FNP in [12].

[11] If the agents are single-minded, then in a PORF mechanism, as long as the prices of larger sets of items are more expensive, an agent's favorite set of items is either the set on which she is single-minded, or the empty set. Thus, we do not need to increase the price of every set by the same amount. As long as we are increasing the prices, an agent's favorite set either remains unchanged, or becomes empty (if the price increase on the set on which she is single-minded is too high). That is, for single-minded agents, we have more flexibility in the process of transforming a feasible mechanism into an FNPW mechanism. Due to space constraint, we do not pursue this further here.

Moreover, the next claim shows that in settings with exactly two substitutable items, the AMD mechanism coincides with MMVIP (but not with VCG).

CLAIM 10. *In settings with exactly two substitutable items, the AMD mechanism coincides with MMVIP.*

PROOF. The proof is by induction on the number of agents. When there is only one agent, this agent faces price 0 for every bundle under the VCG mechanism. This already satisfies S-NSAW, so by Claim 8, we do not need to increase any price in the AMD process. Therefore, when $n = 1$, the AMD mechanism allocates all the items to the only agent for free. The MMVIP mechanism does the same. Hence, when $n = 1$, the AMD mechanism coincides with MMVIP. For the induction step, we assume that the two mechanisms coincide when $n \leq k$. When $n = k + 1$, the price function of the VCG mechanism is defined as: $\chi(\{A\}, O) = v^*_{AB} - v^*_B$, $\chi(\{B\}, O) = v^*_{AB} - v^*_A$, and $\chi(\{AB\}, O) = v^*_{AB}$. Here, A and B are the two items. v^*_A is the highest valuation for A by the agents in O. v^*_B is the highest valuation for B by the agents in O. v^*_{AB} is the highest combined valuation for $\{A, B\}$ by the agents in O (which may be obtained by splitting the items across two different agents, or giving both to the same agent). Since the items are substitutable, $v^*_{AB} \leq v^*_A + v^*_B$. Equivalently, $\chi(\{A\}, O) + \chi(\{B\}, O) \leq \chi(\{AB\}, O)$. Therefore, in the AMD technique, the price of every bundle has to increase by at least $\chi(\{A, B\}, 0) - \chi(\{A\}, O) - \chi(\{B\}, O)$. That is, under the AMD mechanism, the price of A is at least v^*_A, the price of B is at least v^*_B, and the price of $\{A, B\}$ is at least $v^*_A + v^*_B$. These prices are high enough to guarantee the PIA condition, because by the induction assumption, the AMD mechanism coincides with MMVIP for $n \leq k$; so, it follows that the AMD technique results in exactly these prices. They coincide with the prices under the MMVIP mechanism. Therefore, by induction, the AMD mechanism coincides with the MMVIP mechanism for any number of agents, when there are exactly two substitutable items. □

It remains an open question whether there are more general settings in which the AMD mechanism and the MMVIP mechanism coincide.

Finally, we compare the revenue and allocative efficiency of the VCG mechanism, the Set mechanism[12], the MMVIP mechanism, and the AMD mechanism. It should be noted that the VCG mechanism is not FNPW in general. We use it as a benchmark.

We consider a combinatorial auction with two items $\{A, B\}$ and five agents $\{1, 2, \ldots, 5\}$.[13] We denote agent i's valuation for set $S \subseteq \{A, B\}$ by v_i^S. We consider two scenarios, one with valuations displaying substitutability, and the other with valuations displaying complementarity. We randomly generate 1000 instances for each scenario.

Valuations with substitutability: The $v_i^{\{A\}}$ and the $v_i^{\{B\}}$ are drawn independently from $U(0, 1)$ (the uniform distribution from 0 to 1). For all i, $v_i^{\{A,B\}}$ is drawn independently from $U(\max\{v_i^{\{A\}}, v_i^{\{B\}}\}, v_i^{\{A\}} + v_i^{\{B\}})$. In this scenario, AMD and MMVIP coincide. They perform better than the Set mechanism, both in terms of revenue and allocative efficiency.

	VCG	Set	AMD	MMVIP
Revenue	1.285	1.002	1.221	1.221
Efficiency	1.668	1.236	1.550	1.550

Valuations with complementarity: The $v_i^{\{A\}}$ and the $v_i^{\{B\}}$ are still drawn independently from $U(0, 1)$. For all i, $v_i^{\{A,B\}}$ is set to be $(v_i^{\{A\}} + v_i^{\{B\}})(1 + x_i)$, where the x_i are also drawn independently from $U(0, 1)$. It turns out that, in this scenario, Set performs better than AMD and MMVIP, both in terms of revenue and allocative efficiency. (MMVIP performs especially poorly when valuations exhibit complementarity, because every item can potentially have a very large marginal value to another agent, leading to prices that are too high.)

	VCG	Set	AMD	MMVIP
Revenue	1.864	1.849	1.288	0.594
Efficiency	2.372	2.365	1.565	0.721

Thus, when there are two items and five agents, among these FNPW mechanisms, it seems that Set is most desirable if it likely that there is significant complementarity, and AMD is most desirable if it is likely that there is substitutability. (We cannot use the VCG mechanism unless we are certain that the type space makes VCG FNPW.)

7. WORST-CASE EFFICIENCY RATIO OF FNPW MECHANISMS

Yokoo *et al.* [19] proved that in general combinatorial auction settings, there exists no efficient FNP mechanisms. [7] further showed that, under a minor condition called IIG (described below), the worst-case efficiency ratio of any feasible FNP mechanism is at most $\frac{2}{m+1}$.[14]

Definition 13. **Independence of irrelevant good (IIG) [7].** Suppose agent i is winning all the items. If we add an additional item that is only wanted by i, then i still wins all the items.

Given the agents' reported types, the efficiency ratio of a mechanism is defined as the ratio between the achieved allocative efficiency and the optimal allocative efficiency (payments are not taken into consideration). The worst-case efficiency ratio of this mechanism is the minimal such ratio over all possible type profiles.

Example 2. The worst-case efficiency ratio of the Set mechanism is at least $\frac{1}{m}$ [7]. Let v be the winning agent's valuation for the grand bundle. The allocative efficiency of the Set mechanism is v. The optimal allocative efficiency is at most mv, since there are at most m winners in the optimal allocation, and a winner's valuation (for the items she won) is at most v.

Our next theorem is that $\frac{1}{m}$ is a strict upper bound on the efficiency ratios of feasible FNPW mechanisms. That is, the Set mechanism is worst-case optimal in terms of efficiency ratio. Of course, this is only a worst-case analysis, which does not preclude FNPW mechanisms from performing well most of the time.

THEOREM 6. *The worst-case efficiency ratio of any feasible FNPW mechanism is at most $\frac{1}{m}$, if IIG holds, even with single-minded bidders.*

[12] The MB mechanism and the Set mechanism coincide in our experimental setup (the whole bundle is every agent's minimal bundle).

[13] We only focused on these tiny auctions because the AMD technique is computationally quite expensive. Nevertheless, even the solutions to tiny auctions can be helpful in conjecturing more general mechanisms.

[14] [7] also introduced the ARP mechanism, whose worst-case efficiency ratio is exactly $\frac{2}{m+1}$. However, the ARP mechanism is only FNP for single-minded agents. Our next result implies that ARP is not FNPW, even with single-minded bidders.

PROOF. Let χ be the price function that corresponds to an FNPW mechanism with optimal worst-case ratio. Since the Set mechanism is FNPW, χ's worst-case efficiency ratio is at least $\frac{1}{m}$. We denote item i by s_i. We consider the following types:

θ_a: the type of an agent that is single-minded on the grand bundle, with value 1.

θ_i ($i = 1, 2, \ldots, m$): the type of an agent that is single-minded on s_i, with value $1 - \epsilon$. Here, ϵ is a small positive number.

Scenario 1: There are two agents. Agent a has type θ_a. Agent 1 has type θ_1.

Scenario 2: There are two agents. Both agents have type θ_1.

Scenario 3: There are $m+1$ agents. Agent a has type θ_a. Agent i has type θ_i for $i = 1, 2, \ldots, m$.

We first prove that in scenario 1, agent a wins. We start with the special case of $m = 1$. If $\chi(\{s_1\}, \{\theta_1\}) > 1 - \epsilon$, then we consider scenario 2. In scenario 2, both agents can not afford the only item. That is, the efficiency ratio is 0. Hence, we must have $\chi(\{s_1\}, \{\theta_1\}) \leq 1 - \epsilon$. That is, in scenario 1, in the case of $m = 1$, agent a must win. The IIG condition implies that this is also true for cases with $m > 1$.

Since agent a is the only winner in scenario 1, we have $\chi(\{s_1\}, \{\theta_a\}) \geq 1 - \epsilon$ (otherwise, agent 1 would win in scenario 1). ϵ can be made arbitrarily close to 0; hence, $\chi(\{s_1\}, \{\theta_a\}) \geq 1$.

Finally, we consider scenario 3. The price agent 1 faces for s_1 is $\chi(\{s_1\}, \{\theta_a\} \cup (\bigcup_{j \neq 1} \{\theta_j\}))$. According to PIA, this price is at least $\chi(\{s_1\}, \{\theta_a\}) = 1$. That is, agent 1 does not win in scenario 3. By symmetry over the items, agent i does not win for all $i = 1, 2, \ldots, m$. The efficiency ratio in this scenario is then at most $\frac{1}{m(1-\epsilon)}$, which goes to $\frac{1}{m}$ as ϵ goes to 0. $\square$

8. CONCLUSION

We studied a 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. Since FNPW is stronger than FNP, this paper also contributes to the research on false-name-proofness in the traditional sense.

9. REFERENCES

[1] S. Bikhchandani, S. Chatterji, R. Lavi, A. Mu'alem, N. Nisan, and A. Sen. Weak monotonicity characterizes deterministic dominant strategy implementation. *Econometrica*, 74(4):1109–1132, 2006.

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

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

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

[5] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions*. MIT Press, 2006.

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

[7] A. Iwasaki, V. Conitzer, M. Guo, T. Todo, Y. Omori, Y. Sakurai, and M. Yokoo. Worst-case efficiency 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. To appear.

[8] R. Lavi, A. Mu'alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, pages 574–583, 2003.

[9] D. Monderer and M. Tennenholtz. Optimal auctions revisited. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 32–37, Madison, WI, USA, July 1998.

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

[11] T. Sandholm. Limitations of the Vickrey auction in computational multiagent systems. In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS)*, pages 299–306, Keihanna Plaza, Kyoto, Japan, Dec. 1996.

[12] 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.

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

[14] 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. Outstanding paper award.

[15] P. Wurman, M. Wellman, and W. Walsh. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *Proceedings of the Second International Conference on Autonomous Agents (AGENTS)*, pages 301–308, Minneapolis/St. Paul, MN, USA, May 1998.

[16] 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.

[17] M. Yokoo, T. Matsutani, and A. Iwasaki. False-name-proof combinatorial auction protocol: Groves mechanism with submodular approximation. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1135–1142, Hakodate, Japan, 2006.

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

[19] 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.

APPENDIX

Characterizing FNP(W) in Social Choice Settings

Throughout the paper, we have only discussed combinatorial auctions. In this section, we focus on FNP(W)[15] in social choice settings (without payments). Specifically, we present a characteriza-

[15] In these settings, it does not matter whether withdrawal is allowed or not.

tion of FNP(W) social choice functions (without payments). A social choice function f is defined as $f : \{\emptyset\} \cup \Theta \cup \Theta^2 \cup \ldots \rightarrow \Omega$, where Θ is the space of all possible types of an agent, and $\{\emptyset\} \cup \Theta \cup \Theta^2 \cup \ldots$ is the space of all possible profiles (since we do not know how many agents there are). Ω is the outcome space. Let agent i's type be θ_i. Let the types of agents other than i be θ_{-i}. i's valuation for outcome $\omega \in \Omega$ is denoted by $v_i(\theta_i, \omega)$.

First, we present the following straightforward characterization of strategy-proof social choice functions.

CLAIM 11. *A social choice function f is strategy-proof if and only if it satisfies the following condition: $\forall i, \theta_i, \theta_{-i}$, we have $f(\theta_i, \theta_{-i}) \in \arg\max_{\theta'_i} v_i(\theta_i, f(\theta'_i, \theta_{-i}))$.*

PROOF. If the above condition is satisfied, then $\forall i, \theta_i, \theta'_i, \theta_{-i}$, we have $v_i(\theta_i, f(\theta_i, \theta_{-i})) \geq v_i(\theta_i, f(\theta'_i, \theta_{-i}))$. That is, reporting truthfully is a dominant strategy.

If reporting truthfully is a dominant strategy, then $\forall i, \theta_i, \theta'_i, \theta_{-i}$, we have $v_i(\theta_i, f(\theta_i, \theta_{-i})) \geq v_i(\theta_i, f(\theta'_i, \theta_{-i}))$. That is, $\forall i, \theta_i, \theta_{-i}$, we have $v_i(\theta_i, f(\theta_i, \theta_{-i})) \geq \max_{\theta'_i} v_i(\theta_i, f(\theta'_i, \theta_{-i}))$, which is equivalent to $f(\theta_i, \theta_{-i}) \in \arg\max_{\theta'_i} v_i(\theta_i, f(\theta'_i, \theta_{-i}))$. $\square$

That is, an agent always receives her most-preferred choice among outcomes that she can attain with some type report. We are now ready to present the characterization of FNP(W) social choice functions.

CLAIM 12. *Suppose that for every outcome $o \in \Omega$, there exists some type $\theta_i \in \Theta$ such that $\{o\} = \arg\max_{o' \in O} u_{\theta_i}(o')$ (each o is the unique most-preferred outcome for some type). Then, a strategy-proof and individually rational social choice function f is FNP(W) if and only if it satisfies the following condition: $\forall i, \theta_{-i}, \theta_0$, we have $\{f(\theta_i, \theta_{-i}) | \theta_i \in \Theta\} \supseteq \{f(\theta_i, \theta_{-i} \cup \{\theta_0\}) | \theta_i \in \Theta\}$. That is, with an additional other agent, the set of outcomes that an agent can choose decreases or stays the same.*

PROOF. We first show that if f is FNP(W), then the condition must be satisfied. Suppose not, that is, for some i, θ_{-i}, θ_0, there exists some $o \in \{f(\theta_i, \theta_{-i} \cup \{\theta_0\}) | \theta_i \in \Theta\} \setminus \{f(\theta_i, \theta_{-i}) | \theta_i \in \Theta\}$. Then, by assumption, there exists some $\theta_i \in \Theta$ such that $\{o\} = \arg\max_{o' \in O} u_{\theta_i}(o')$. It follows that an agent facing type profile θ_{-i} cannot obtain o with a single report, but can obtain it by reporting both θ_0 and some other type (such as, by strategy-proofness, θ_i). Because o is her unique most-preferred outcome, she prefers to engage in this manipulation, contradicting FNP(W).

Conversely, we show that if the condition is satisfied, then f is FNP(W). By assumption, f is strategy-proof and individually rational, so we only need to check that an agent has no incentive to use multiple identifiers. Suppose that o is an outcome that i can obtain when facing θ_{-i} by submitting multiple identities. Because the set of choices is nonincreasing in the number of identifiers used according to the condition, it must be that $o \in \{f(\theta_i, \theta_{-i}) | \theta_i \in \Theta\}$. Hence, there is no reason for her to use more than one identity. $\square$