Worst-case efficiency ratio in false-name-proof combinatorial auction mechanisms

Atsushi Iwasaki[1], Vincent Conitzer[2], Yoshifusa Omori[1], Yuko Sakurai[1], Taiki Todo[1], Mingyu Guo[2], and Makoto Yokoo[1]
[1]Kyushu University, Fukuoka, 819-0395, Japan
{iwasaki@, omori@agent., sakurai@agent.,todo@agent., yokoo@}is.kyushu-u.ac.jp
[2]Duke University, Durham, NC 27708, USA
{conitzer, mingyu}@cs.duke.edu

ABSTRACT

This paper analyzes the worst-case efficiency ratio of false-name-proof combinatorial auction mechanisms. False-name-proofness generalizes strategy-proofness by assuming that a bidder can submit multiple bids under fictitious identifiers. Even the well-known Vickrey-Clarke-Groves mechanism is not false-name-proof. It has previously been shown that there is no false-name-proof mechanism that always achieves a Pareto efficient allocation. Consequently, if false-name bids are possible, we need to sacrifice efficiency to some extent. This leaves the natural question of how much surplus must be sacrificed. To answer this question, this paper focuses on worst-case analysis. Specifically, we consider the fraction of the Pareto efficient surplus that we obtain and try to maximize this fraction in the worst-case, under the constraint of false-name-proofness. As far as we are aware, this is the first attempt to examine the worst-case efficiency of false-name-proof mechanisms.

We show that the worst-case efficiency ratio of any false-name-proof mechanism that satisfies some apparently minor assumptions is at most $2/(m+1)$ for auctions with m different goods. We also observe that the worst-case efficiency ratio of existing false-name-proof mechanisms is generally $1/m$ or 0. Finally, we propose a novel mechanism, called the *adaptive reserve price* mechanism that is false-name-proof when all bidders are single-minded. The worst-case efficiency ratio is $2/(m+1)$, i.e., optimal.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multi-agent systems*; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Theory, Economics, Design

Keywords

Mechanism design, Combinatorial auctions, Worst-case analysis, False-name-proofness

Cite as: Worst-case efficien y ratio in false-name-proof combinatorial auction mechanisms, Atsushi Iwasaki, Vincent Conitzer, Yoshifusa Omori, Yuko Sakurai, Taiki Todo, Mingyu Guo, and Makoto Yokoo, *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. 633-640

1. INTRODUCTION

In a combinatorial auction, multiple goods are simultaneously for sale, and, in general, bidders can express arbitrary valuation functions over subsets of the goods. This allows bidders to express substitutability and complementarity of goods in their valuations. The recent book by Cramton *et al.* [3] gives a thorough survey of the theory and practice of combinatorial auctions.

One desirable characteristic of an auction mechanism is that it is *strategy-proof*. A mechanism is strategy-proof if, for each bidder, declaring his/her true valuation is a *dominant strategy*, i.e., an optimal strategy regardless of the actions of other bidders. The revelation principle [12] states that we can, without loss of generality, restrict our attention to strategy-proof mechanisms if we require implementation in dominant strategies. In other words, if a certain property (e.g., Pareto efficiency) can be achieved using some auction mechanism in a dominant-strategy equilibrium, i.e., by a combination of dominant strategies for the bidders, then the property can also be achieved using a strategy-proof auction mechanism.

It can be argued that using a strategy-proof mechanism is especially advantageous in Internet auctions, where we need to worry about the privacy of the bids. For example, if we use the first-price sealed-bid auction, which is not strategy-proof, then the bids must be securely concealed until the auction ends. On the other hand, if we use a strategy-proof mechanism, then an individual bidder does not care to know the others' bids; consequently, such security issues become less critical. The Vickrey-Clarke-Groves (VCG) mechanism is a strategy-proof mechanism that can be applied to combinatorial auctions (resulting in what is also known as the Generalized Vickrey Auction). We say that an auction mechanism is *Pareto efficient* when the sum of all participants' utilities (including that of the auctioneer, who receives the payments)–i.e., the social surplus–is maximized in a dominant-strategy equilibrium. The VCG mechanism always satisfies Pareto efficiency.

However, declaring valuations untruthfully is only one way to manipulate the mechanism. Another way is for one bidder to pretend to be multiple bidders. Such *false-name bids* [17] are especially feasible in Internet auctions due to their relative anonymity. False-name bids are bids submitted under fictitious names, e.g., multiple e-mail addresses. This type of manipulation is very difficult to detect, since identifying each participant on the Internet is virtually impossible.

We say a mechanism is *false-name-proof* if, for each bidder, declaring his true valuation function using a single identifier is a dominant strategy (even though the bidder can choose to use multiple identifiers). However, the VCG mechanism is not false-name-proof [17]. Hence, it cannot be used to achieve a Pareto efficient allocation if false-name bids are possible. In fact, *no* false-name-proof mechanism satisfies Pareto efficiency [17]. Therefore, we need to sacrifice efficiency to some extent when false-name bids are possible. This leaves the natural question of *how much* surplus must be sacrificed. As far as we know, there has not yet been any theoretical analysis of this question.

There are several ways to proceed. One is to take a Bayesian perspective and construct a mechanism that maximizes *expected* surplus based on the prior distribution from which bidders' private values are drawn. Indeed, combinatorial auctions and similar problems have recently been studied from the Bayesian perspective [6, 4, 10]. However, in this paper, we take a prior-free approach: we focus on worst-case analysis. Specifically, we consider the fraction of the Pareto efficient surplus that we obtain and then we try to maximize this fraction in the worst-case, under the constraint of false-name-proofness.

Worst-case analysis is commonly used in the recent mechanism design literature, especially by computer scientists. Several recent papers focus on maximizing the profit of an auction according to a worst-case competitive analysis (see [13] as an extensive survey). Other work tries to redistribute as much revenue back to the bidders as possible according to a worst-case criterion [7, 11]. The loss of efficiency in network games due to selfish user behavior has been studied in terms of the "price of anarchy [14]" and the "price of stability [1]." Furthermore, Archer, Tardos, Talwar and others [2, 8] study a hiring-a-team problem. Significant insight can be gained from an understanding of worst-case performance. It allows an uninformed or partially informed auctioneer to evaluate the trade-off between an auction based on assumptions about the distribution of bidder valuations (which may or may not be correct), and an auction designed to work as well as possible under unknown and worst-case market conditions.

Let us briefly describe the organization and the main contributions of this paper. First, in Section 2, we formalize combinatorial auctions according to a price-based description called the Price-Oriented, Rationing-Free mechanism [15]. Second, in Section 3, under a mild and reasonable *independence of irrelevant good* conditions, we show that the worst-case efficiency ratio of any false-name-proof mechanism is at most $\frac{2}{m+1}$. This upper bound holds even if we assume that all bidders are single-minded.

Third, in Section 4, we examine the worst-case efficiency ratio of existing false-name-proof mechanisms and show that the worst-case efficiency ratio of a trivial mechanism, called the *Set* mechanism, is $\frac{1}{m}$, while it is 0 for other more sophisticated mechanisms. If we assume that all bidders are single-minded, then a mechanism called the *Minimal Bundle (MB)* mechanism also achieves the worst-case efficiency ratio $1/m$, and it dominates the Set mechanism in terms of efficiency.

Finally, in Section 5, we develop a new mechanism called the *adaptive reserve price (ARP)* mechanism, which is false-name-proof when all bidders are single-minded, and its worst-case efficiency ratio is $\frac{2}{m+1}$. This ratio matches the theoretical upper bound.

2. MODEL

Let $N = \{0, 1, 2, \ldots, n\}$ be the set of bidders and let $M = \{g_1, g_2, \ldots, g_m\}$ be the set of goods. Each bidder i has preferences over the bundles $B \subseteq M$. We model this by supposing that bidder i privately observes a parameter, or signal, θ_i, which determines his/her preferences. We refer to θ_i as the *type* of bidder i, which is drawn from Θ. We also assume a *quasi-linear, private value* model with *no allocative externalities*, i.e, the utility of bidder i when i obtains bundle (subset of goods) $B \subseteq M$ and pays p is assumed to equal $v(B, \theta_i) - p$. We assume that $v(\emptyset, \theta_i) = 0$ and that there is *free disposal*, i.e., $v(B', \theta_i) \geq v(B, \theta_i)$ for all $B' \supseteq B$.

In a context where false-name bids are not possible, an auction mechanism is (dominant-strategy) *incentive compatible* (or *strategy-proof*) if declaring the true type/valuation is a dominant strategy for each bidder, i.e., an optimal strategy regardless of the actions of other bidders.

In this paper, we extend the traditional definition of incentive compatibility so that it can address false-name manipulations, i.e., we say that an auction mechanism is (dominant strategy) incentive compatible if using a single identifier and declaring the true type under that identifier is a dominant strategy for each bidder. To distinguish between the traditional and extended definitions of incentive compatibility, we refer to the traditional concept as strategy-proofness and to the extended definition as *false-name-proofness*.

We also restrict our attention to individually rational mechanisms, where no participant suffers any loss in a dominant-strategy equilibrium, i.e., the payment never exceeds the evaluation value of the obtained goods. Moreover, we restrict our attention to deterministic mechanisms, which always obtain the same outcome for the same input.

Under these assumptions, we describe combinatorial auction mechanisms according to a general framework for describing strategy-proof mechanisms called Price-Oriented, Rationing-Free (PORF) mechanisms [15]. By describing a mechanism as a PORF, proving that the mechanism is strategy-proof or false-name-proof becomes much easier. Similar price-based representations have also been presented by others, including [9].

Under a PORF mechanism, each bidder i declares his type $\tilde{\theta}_i$, which is not necessarily the true type θ_i. For each i and bundle $B \subseteq M$, the price $p(B, \tilde{\Theta}_{-i})$ is defined, where $\tilde{\Theta}_{-i}$ is a set of declared types other than i. This price must be determined independently of i's declared type $\tilde{\theta}_i$, while it can be dependent on declared types of other bidders. We assume the pricing rule is *anonymous*, i.e., a single pricing rule is used for all agents and it is defined on the set of other agents' types. Thus, if two agents have exactly the same type, they will face the same prices. We assume $p(\emptyset, \tilde{\Theta}_{-i}) = 0$ and for any $B \subseteq B'$, $p(B, \tilde{\Theta}_{-i}) \leq p(B', \tilde{\Theta}_{-i})$ holds.

A PORF mechanism allocates to bidder i a bundle B^* so that $B^* = \arg\max_{B \subseteq M} v(B, \tilde{\theta}_i) - p(B, \tilde{\Theta}_{-i})$. Bidder i pays $p(B^*, \tilde{\Theta}_{-i})$. If there exist multiple bundles that maximize i's utility, one of these bundles is allocated.

The pricing rule must be defined so that the allocation satisfies *allocation feasibility*, i.e., for two bidders i, j and bundles allocated to these bidders B_i^* and B_j^*, $B_i^* \cap B_j^* = \emptyset$ holds. This condition guarantees that a PORF mechanism is strategy-proof. The price of bidder i for each possible bundle is determined independently of i's declared type, and he/she can obtain the bundle that maximizes his/her utility independently of the allocations of other bidders, i.e., the mechanism is rationing-free.

Furthermore, we are introducing *symmetry for prices across goods* to the pricing rule. Let us define ρ as a permutation of goods or types. $\rho(B)$ is a bundle where the goods in B are renamed. $\rho(\theta_i)$ is a new type θ'_i, where $v(B, \theta_i) = v(\rho(B), \theta'_i)$ holds for all $B \subseteq M$. $\rho(\tilde{\Theta})$ is a set $\{\rho(\theta_i) | \theta_i \in \tilde{\Theta}\}$. Now, we are ready to introduce symmetry for prices across goods.

Definition 1 *A pricing rule is symmetric across goods if, for all i, ρ, $\tilde{\Theta}$, B, $p(B, \tilde{\Theta}_{-i}) = p(\rho(B), \rho(\tilde{\Theta}_{-i}))$.*

This condition means that the names of goods do not affect the prices of goods. Almost all well-known mechanisms, in particular all existing false-name-proof mechanisms, satisfy this condition.

Finally, we introduce the worst-case efficiency ratio. Let $\tilde{\Theta}$ be a (multi)set of declared types (bids). Let $s^{\mathcal{M}}(\tilde{\Theta})$ be the social surplus that mechanism $\mathcal{M}$ achieves on this set of bids, and let $s^*(\tilde{\Theta})$ be the corresponding Pareto efficient social surplus. The ratio of $\mathcal{M}$ is defined as:

$$\inf_{\tilde{\Theta}} \frac{s^{\mathcal{M}}(\tilde{\Theta})}{s^*(\tilde{\Theta})}.$$

3. UPPER BOUND ON THE WORST-CASE EFFICIENCY RATIO

This section prove that the worst-case efficiency ratio of any false-name-proof mechanism is at most $\frac{2}{m+1}$ under an apparently minor condition in false-name-proof combinatorial auction mechanisms with m different goods.

Definition 2 (Independence of irrelevant good (IIG)) *Assume bidder i is winning all goods. If we add an additional good that is wanted only by bidder i, and the valuation of i for all goods is larger than or equal to some constant c, then i still wins all goods.*

The condition is intuitively reasonable. Accordingly, it is satisfied in almost all well-known mechanisms, in particular in all existing false-name-proof mechanisms, as far as the authors aware. This is true for a mechanism that uses predefined reserve prices, such as the LDS mechanism [16] assuming c is large enough compared to the reserve price.

It should be emphasized that the IIG is different from typical Independence of Irrelevant Alternatives (IIA), which are often quite strong and apply to a wide variety of situations. This condition applies only to very specific situations, and we feel that it is quite mild. Nevertheless, while introducing some technical conditions to characterize a certain class of mechanisms is a common practice in mechanism design, we hope to remove this condition in our future work.

Consider the following situation (Case 1), with m goods and two bidders with the single-minded valuation functions.

Case 1

bidder 0: c for goods g_1 to g_m (all or nothing),
bidder 1: $c - \epsilon$ for good g_1.

Lemma 1 *Any deterministic, symmetric, strategy-proof mechanism satisfying IIG whose worst-case efficiency ratio is non-zero will allocate all the goods to bidder 0 in Case 1.*

Proof. To prove that the lemma holds for all m, we use mathematical induction on m.

Base case.

Case 2

bidder 0: c for good g_1,
bidder 1: $c - \epsilon$ for good g_1.

Let us prove that bidder 0 wins in Case 2. To derive a contradiction, let us first assume that the mechanism allocates good g_1 to bidder 1 in Case 2. Consider the following case:

Case 3

bidder 0: c for good g_1,
bidder 1: c for good g_1.

Since we assume that the mechanism is deterministic, either bidder 0 or 1 must lose. Then, the losing bidder can under-declare his valuation and make the situation identical to Case 2. This bidder will win the good for at most $c - \epsilon$ and hence have positive utility. This contradicts the assumption that the mechanism is strategy-proof. Thus, bidder 1 cannot win in a strategy-proof mechanism. Also, since we assume that the worst-case efficiency ratio is non-zero, bidder 0 must win in Case 2. Therefore, the lemma holds for $m = 1$.

Induction step.

Let us assume that the lemma holds for $m = k$, i.e., the mechanism allocates all goods to bidder 0 in this case. We show that it still allocates all goods to bidder 0 in the case where $m = k + 1$:

Case 4

bidder 0: c for goods g_1 to g_{k+1},
bidder 1: $c - \epsilon$ for good g_1.

By IIG and the induction hypothesis, bidder 0 still wins all goods, since only bidder 0 is interested in good $k + 1$. □

Now, consider Case 5 with m goods and $n + 1 = m + 1$ bidders.

Case 5

bidder 0: c for goods g_1 to g_m,
bidder 1: $c - \epsilon$ for good g_1,
bidder j ($2 \le j \le m$): $c/2 - \epsilon$ for good g_j.

Lemma 2 *Any deterministic, symmetric, false-name-proof mechanism satisfying IIG whose worst-case efficiency ratio is non-zero will allocate all the goods to bidder 0 in Case 5.*

Proof. To prove this lemma, we use mathematical induction on the number of bidders, while keeping m fixed, that is, we consider the case where some of the last bidders are removed, without removing the corresponding goods.

Base case.

Case 6

bidder 0: c for goods g_1 to g_m,
bidder 1: $c - \epsilon$ for good g_1,
bidder 2: $c/2 - \epsilon$ for good g_2.

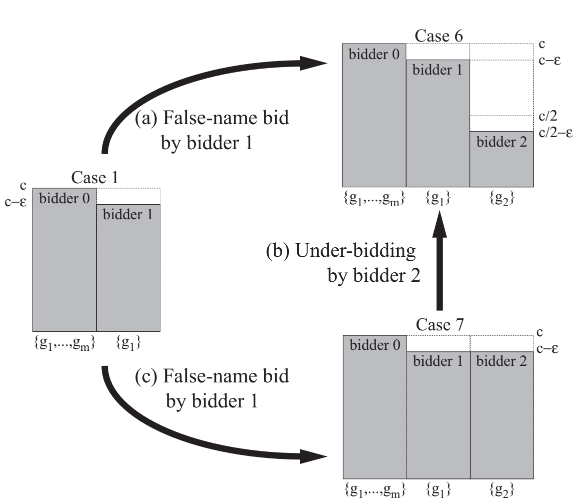

Figure 1: Bids in Case 1, 6, and 7.

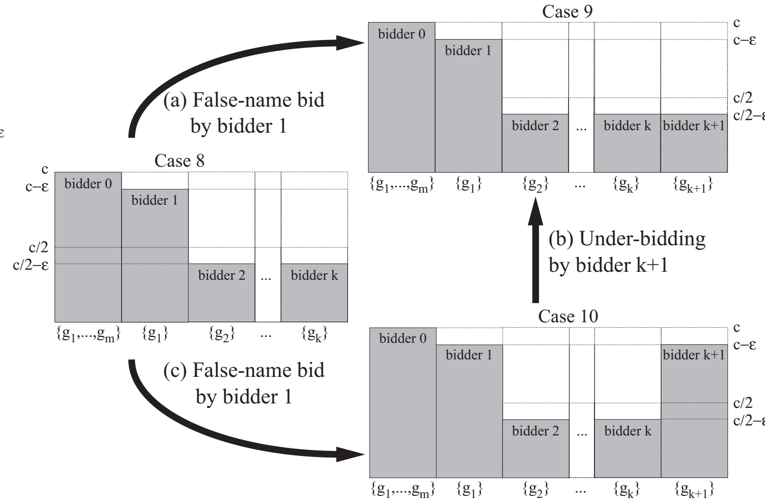

Figure 2: Bids in Case 8, 9, and 10.

First, let us assume that bidder 1 wins and bidder 2 loses. Since false-name-proofness includes strategy-proofness, Lemma 1 holds. Then, in Case 1 where bidder 1 is losing, bidder 1 has an incentive to use a false-name and make the situation identical to Case 6, i.e., bidder 1, who obtains no good in Case 1, can obtain good g_1 for at most $c - \epsilon$ in this way (Figure 1 (a)). This contradicts the assumption that the mechanism is false-name-proof.

Now, let us assume that bidder 2 wins in Case 6. Consider the following case:

Case 7

bidder 0:	c for goods g_1 to g_m,
bidder 1:	$c - \epsilon$ for good g_1,
bidder 2:	$c - \epsilon$ for good g_2.

In Case 7, bidder 2 must win at a price of at most $c/2 - \epsilon$; otherwise, he has an incentive to under-bid to $c/2 - \epsilon$, making the situation identical to Case 6 (Figure 1 (b)). Let us consider permutation ρ, which exchanges the names of g_1 and g_2 while other names remain unchanged. In this situation, $\rho(\theta_2) = \theta_1$, $\rho(\theta_0) = \theta_0$, and $\rho(\{g_1\}) = \{g_2\}$ hold. From the symmetric pricing rule across goods in Definition 1,

$$\begin{aligned} p(\{g_1\}, \{\theta_0, \theta_2\}) = &\; p(\rho(\{g_1\}), \rho(\{\theta_0, \theta_2\})) \\ = &\; p(\{g_2\}, \{\theta_0, \theta_1\}) \leq c/2 - \epsilon. \end{aligned}$$

Thus, bidder 1 also wins and his payment is at most $c/2 - \epsilon$.

But then, in Case 1, where bidder 1 is losing from Lemma 1 by false-name bidding, he can make the situation identical to Case 7, thereby winning goods g_1 and g_2 and paying at most $2(c/2 - \epsilon)$ (Figure 1 (c)). Thus, for bidder 1, using a false-name manipulation is profitable, since his utility increases from 0 to at least $c - \epsilon - 2(c/2 - \epsilon) = \epsilon$. This contradicts the assumption that the mechanism is false-name-proof.

Thus, neither bidder 1 nor bidder 2 is a winner in Case 6. Since we assume the worst-case efficiency ratio is non-zero, bidder 0 must win. Therefore, the lemma holds for $n = 2$.

Induction step.

Let us assume that the lemma holds for $n = k$, i.e., the mechanism allocates m goods to bidder 0 in Case 8.

Case 8

bidder 0:	c for goods g_1 to g_m,
bidder 1:	$c - \epsilon$ for good g_1,
bidder j $(2 \leq j \leq k)$:	$c/2 - \epsilon$ for good g_j.

Then, we show that it still allocates m goods to bidder 0 in the following case with $n = k + 1$ (Case 9).

Case 9

bidder 0:	c for goods g_1 to g_m,
bidder 1:	$c - \epsilon$ for good g_1,
bidder j $(2 \leq j \leq k+1)$:	$c/2 - \epsilon$ for good g_j.

First, assume that bidder 1 wins, but none of the bidders from 2 to $k + 1$ win. Then, in Case 8, bidder 1 has an incentive to use a false-name-bid and make the situation identical to Case 9 (Figure 2 (a)). Thus, it is impossible for the mechanism to choose only bidder 1 as the winner. Second, if bidder 0 never becomes the winner, from the non-zero efficiency assumption, at least one of the bidders from 2 to k+1 wins. Thus, w.l.o.g., we assume that bidder $k + 1$ wins. Then, consider the following case:

Case 10

bidder 0:	c for goods g_1 to g_m,
bidder 1:	$c - \epsilon$ for good g_1,
bidder j $(2 \leq j \leq k)$:	$c/2 - \epsilon$ for good g_j,
bidder $k + 1$:	$c - \epsilon$ for good g_{k+1}.

In Case 10, bidder $k+1$ must win and his payment must be at most $c/2 - \epsilon$; otherwise, he has an incentive to under-bid to $c/2 - \epsilon$, making the situation identical to Case 9 (Figure 2 (b)). Let us consider permutation ρ, which exchanges the names of g_1 and g_{k+1} while other names remain unchanged. Then, in this situation, $\rho(\theta_{k+1}) = \theta_1$, $\rho(\theta_i) = \theta_i$ for $i = 0, 2, \ldots, k$, and $\rho(\{g_1\}) = \{g_{k+1}\}$ hold. From the symmetric pricing rule across goods in Definition 1,

$$\begin{aligned} &p(\{g_1\}, \{\theta_0, \theta_2, \ldots, \theta_{k+1}\}) = p(\rho(\{g_1\}), \rho(\{\theta_0, \theta_2, \ldots, \theta_{k+1}\})) \\ &= p(\{g_{k+1}\}, \{\theta_0, \theta_1, \ldots, \theta_k\}) \leq c/2 - \epsilon. \end{aligned}$$

Thus, bidder 1 also wins and his payment is at most $c/2 - \epsilon$.

However, in Case 8, bidder 1 can submit a false-name bid as bidder $k+1$ and make the situation identical to Case 10 (Figure 2 (c)). Then, bidder 1 obtains goods g_1 and g_{k+1}, pays at most $2(c/2 - \epsilon)$, and his utility is at least $c - \epsilon - 2(c/2 - \epsilon) = \epsilon > 0$. Thus, using a false-name-bid is profitable for bidder 1. This contradicts the assumption that the mechanism is false-name-proof. It follows that none of the bidders from 2 to $k+1$ can win in Case 9. Since the mechanism has a non-zero worst-case efficiency ratio, bidder 0 must win. □

Theorem 1 *For any deterministic, symmetric, false-name-proof mechanism satisfying IIG with m goods and* $n+1$ ($\geq m+1$) *bidders, the worst-case efficiency ratio is at most* $\frac{2}{m+1}$. *This is true even if we assume all bidders are single-minded.*

Proof. From Lemma 2, in Case 5, any such mechanism allocates all of the goods to bidder 0 unless its ratio is 0 (the proof of Lemma 2 only uses single-minded bids). Thus, the social surplus is c, while the Pareto efficient social surplus is $(c - \epsilon) + (m-1)(c/2 - \epsilon)$. As a result, the worst-case efficiency ratio is at most $\frac{2}{m+1}$. □

4. EXISTING FALSE-NAME-PROOF MECHANISMS

This section examines the worst-case efficiency ratios of three existing false-name-proof mechanisms. First, the Set mechanism is one of the simplest false-name-proof ones. It allocates all goods to a single bidder, namely, the bidder with the largest valuation for the grand bundle of all goods. Effectively, it sells the grand bundle as a single good, in a Vickrey/second-price auction.

Theorem 2 *For Set with m different goods, the worst-case efficiency ratio is* $\frac{1}{m}$. *This is true even if we assume all bidders are single-minded.*[1]

Proof. First, we show that the worst-case efficiency ratio is at least $\frac{1}{m}$. Let v^{max} be the highest valuation for the grand bundle (and hence, by free disposal, for any bundle); this is the efficiency obtained by Set. In any allocation, there can be at most m winners, and each winner can have a valuation of at most v^{max}, hence no allocation has efficiency greater than mv^{max}.

Let us prove that the ratio is at most $\frac{1}{m}$. Suppose that bidder 0 values c on the bundle of all goods, whereas bidder $j \in [1, m]$ value $c - \epsilon$ on only g_j. Set allocates all goods to bidder 0, obtaining an efficiency of c even though $m(c - \epsilon)$ is possible. Thus, the ratio is arbitrarily close to $\frac{1}{m}$. This is true even if we assume all bidders are single-minded. □

Second, the Minimal Bundle (MB) [15] mechanism can be thought of as a modified version of Set. To describe this mechanism, let us first introduce a *minimal bundle.* A bundle B is called *minimal* for bidder i if for all $B' \subsetneq B$, $v(B', \theta_i) < v(B, \theta_i)$ holds. Instead of allocating all goods M to bidder i who has the highest valuation, we first allocate $B \subseteq M$ to i, where B is a minimal bundle of i. Then, we allocate $B' \subseteq M \setminus B$ to another bidder j who has the highest remaining valuation, where B' is a minimal bundle of j, and so on. The payment for an allocated bundle B is equal to the highest valuation of another bidder for a bundle that is minimal and conflicting with B. For a single-minded bidder, the minimal bundle for him/her is uniquely determined.

[1]Dobzinski and Nisan [5] discuss maximal-in-range mechanisms, which include Set as a special case, though it does not directly discuss the worst-case efficiency ratio for this special case.

Theorem 3 *For MB with m different goods, the worst-case efficiency ratio is* $\frac{1}{m}$ *if we assume that all bidders are single-minded. In general, the worst-case efficiency ratio is 0.*

Proof. If all bidders are single-minded, then the obtained social surplus of MB always dominates that of Set. This is because the winner of Set is always included in the winners of MB. Thus, the ratio is at least $1/m$. Also, in the situation of Theorem 2, the outcomes of Set and MB are identical. Thus, the ratio is $\frac{1}{m}$.

For the general case, i.e., bidders who are not necessarily single-minded, suppose that bidder 1 values 1 on g_1, 2ϵ on g_2, and 1 on the both goods, whereas bidder 2 values $1 - \epsilon$ on only g_1. The Pareto efficient allocation gives g_2 to bidder 1 and g_1 to bidder 2; the social surplus is $1 + \epsilon$. MB can give bidder 1 either g_1 or g_2. The prices for g_1 and g_2 are $1 - \epsilon$ and 0, respectively. For bidder 1, obtaining g_2 is better and g_1 remains unsold, and the social surplus is 2ϵ. Thus, the ratio is $2\epsilon/(1 + \epsilon)$, which is arbitrarily close to 0. □

Finally, let us examine the Leveled Division Set (LDS) mechanism [16]. Since it is quite complicated and we focus on the worst-case efficiency ratio, let us just note that it must predetermine the reserve prices for each single good. For two goods g_1 and g_2, if both $r_{\{g_1\}}$ and $r_{\{g_2\}}$ are 0, LDS is identical to Set. So, we assume either $r_{\{g_1\}}$ or $r_{\{g_2\}}$ is non-zero and derive the worst-case efficiency ratio.

Theorem 4 *For LDS with m different goods, where at least one reserve price is non-zero, the worst-case efficiency ratio is zero. This is true even if we assume all bidders are single-minded.*

Proof. Let us assume, w.l.o.g., that the reserve price $r_{\{g_1\}}$ for good g_1 is positive. If there exists only one bidder i, whose valuation for g_1 is $r_{\{g_1\}} - \epsilon$, then LDS allocates no good, resulting in a social surplus of 0, while the Pareto efficient allocation allocates g_1 to bidder i, resulting in a social surplus of $r_{\{g_1\}} - \epsilon$. Thus, the ratio is 0. This is true even if we assume all bidders are single-minded. □

5. AN ADAPTIVE RESERVE PRICE MECHANISM

This section introduces a new mechanism that we call the Adaptive Reserve Price (ARP) mechanism for single-minded bidders.

Consider an auction with m goods, $g_1, \ldots, g_m$. Let us assume that all bidders are single-minded. A bidder is allowed to bid on either the bundle of all goods or an individual good. If a bidder bids on a bundle of at least two goods, it is treated as a bid on all goods. First, let us precisely define the ARP mechanism as a PORF mechanism. Let us denote the highest bids for each bundle/good other than bidder i as $v^{-i}_{\{g_1,\ldots,g_m\}}$, $v^{-i}_{\{g_1\}}$, $\ldots$, $v^{-i}_{\{g_m\}}$. Then, let us denote the good that has the k-th highest bid among $g_1, \ldots, g_m$ as $g_{(k)}$, i.e., we assume $v^{-i}_{\{g_{(1)}\}} \geq \ldots \geq v^{-i}_{\{g_{(m)}\}}$ holds.

ARP first determines prices of bundles as follows:

$$p(\{g_1,\ldots,g_m\},\Theta_{-i}) = \max(v^{-i}_{\{g_1,\ldots,g_m\}}, v^{-i}_{\{g_{(1)}\}}, 2v^{-i}_{\{g_{(2)}\}}).$$

$$p(\{g_{(1)}\},\Theta_{-i}) = \begin{cases} \max(v^{-i}_{\{g_{(1)}\}}, v^{-i}_{\{g_1,\ldots,g_m\}}/2) \\ \text{if } v^{-i}_{\{g_1,\ldots,g_m\}} < 2v^{-i}_{\{g_{(2)}\}}, \\ \max(v^{-i}_{\{g_{(1)}\}}, v^{-i}_{\{g_1,\ldots,g_m\}}) \\ \text{otherwise.} \end{cases}$$

$$\forall k \in [2,m], p(\{g_{(k)}\},\Theta_{-i}) = \begin{cases} \max(v^{-i}_{\{g_{(2)}\}}, v^{-i}_{\{g_1,\ldots,g_m\}}/2) \\ \text{if } v^{-i}_{\{g_1,\ldots,g_m\}} < 2v^{-i}_{\{g_{(1)}\}}, \\ \max(v^{-i}_{\{g_{(1)}\}}, v^{-i}_{\{g_1,\ldots,g_m\}}) \\ \text{otherwise.} \end{cases}$$

Then, for bidder i, a bundle B^* is allocated, which is defined by $B^* = \arg\max_{B\subseteq M} v(B,\tilde{\theta}_i) - p(B,\tilde{\Theta}_{-i})$. Bidder i obtaining bundle B^* pays $p(B^*,\tilde{\Theta}_{-i})$.

If there exist multiple bundles that maximize i's utility, ARP allocates one of these bundles. Since each agent is single-minded, this occurs only when an agent is indifferent between obtaining his desired bundle and obtaining nothing (an empty bundle). If this is the case, we assume ARP chooses one possible allocation that satisfies allocation feasibility. In particular, whenever an allocation for two bidders is possible, ARP chooses that allocation, i.e., there are at most two winners.

Next, we show a more informal description of the ARP mechanism. Let us denote the highest bids for each bundle/good as $v^*_{\{g_1,\ldots,g_m\}}$, $v^*_{\{g_1\}}$, $\ldots$, $v^*_{\{g_m\}}$ and w.l.o.g., assume $v^*_{\{g_1\}} \geq, \ldots, \geq v^*_{\{g_m\}}$. Also, let us assume $v^*_{\{g_1,\ldots,g_m\}}$, $v^*_{\{g_1\}}$, $\ldots$, $v^*_{\{g_m\}}$ are submitted by bidder $0,1,\ldots,m$, respectively. Then, in ARP, if $v^*_{\{g_1,\ldots,g_m\}} \geq 2v^*_{\{g_2\}}$, all goods are allocated to bidder 0. Otherwise, g_1 is allocated to bidder 1 and g_2 is allocated to bidder 2. No other goods are then allocated.

Now, let us describe the basic idea of ARP. Set and MB choose as winners the bidders with the highest bids on bundles without considering how many goods are in those bundles. In these, the worst-case efficiency ratio is no more than $1/m$, as shown in Theorems 2 and 3. On the other hand, LDS can choose a bidder who demands a small bundle, even if he does not submit the highest bid, by using reserve prices. However, because the reserve prices are fixed, the ratio is zero, as shown in Theorem 4. The basic idea of ARP is to base the reserve prices on the (other bidders') declared bids. The reserve price on the set of all goods is $2v^*_{\{g_2\}}$, which is obtained by doubling the second highest bid among ones for each single good. The reserve prices on good g_1 and g_2 are $v^*_{\{g_1,\ldots,g_m\}}/2$.

The following three examples illustrate how ARP works.

Example 1 Let us consider an auction with three bidders, 0, 1, and 2, and two goods, g_1 and g_2. The bidders have the following valuations:

bidder 0:	5 for goods $\{g_1, g_2\}$,
bidder 1:	4 for good g_1,
bidder 2:	3 for good g_2.

In VCG, bidder 1 and 2 buy g_1 and g_2 at 2 and 1, while in MB, bidder 0 buys both goods at 4. In ARP, the price for bidder 0 to buy both goods is

$$\max(v^{-0}_{\{g_1,\ldots,g_m\}}, v^{-0}_{\{g_{(1)}\}}, 2v^{-0}_{\{g_{(2)}\}}) = \max(0,4,6) = 6.$$

Since his utility when obtaining all goods is negative, he buys nothing. On the other hand, since $v^{-1}_{\{g_1,\ldots,g_m\}} = 5$ is smaller than $2v^{-1}_{\{g_{(2)}\}} = 6$, the price for bidder 1 to buy g_1 is $\max(v^{-1}_{\{g_1\}}, v^{-1}_{\{g_1,\ldots,g_m\}}/2) = \max(0,2.5) = 2.5$.

Since his utility when obtaining g_1 is positive, he buys g_1 at 2.5. Similarly, bidder 2 buys g_2 at 2.5. In this case, ARP achieves the Pareto efficient allocation. Even if bidder 1 and 2 are false-name bids by a single bidder whose valuation for both goods is 7, using false-name bids does not increase his utility, since the total payment does not change.

Example 2 There are three bidders and two goods.

bidder 0:	5 for goods $\{g_1, g_2\}$,
bidder 1:	4 for good g_1,
bidder 2:	2 for good g_2.

In VCG and MB, the winners are the same as in Example 1. In ARP, the price for bidder 0 to buy both goods is $\max(0,4,4) = 4$. Since his utility when obtaining them is positive, he buys both goods. In VCG, if bidder 1 and 2 are false-name bids of a single bidder whose valuation for both goods is 6, submitting false-name bids benefits that bidder because his payment decreases from 5 to $3+1=4$. In ARP, submitting false-name bids does not help the bidder because he does not win any good.

Example 3 There are three bidders and two goods.

bidder 0:	5 for goods $\{g_1, g_2\}$,
bidder 1:	6 for good g_1,
bidder 2:	2 for good g_2.

In VCG, bidder 1 and 2 buy g_1 at 3 and g_2 at 0, while in MB, bidder 1 buys g_1 at 5. In ARP, since $v^{-1}_{\{g_1,\ldots,g_m\}} = 5$ is greater than $2v^{-1}_{\{g_{(2)}\}} = 4$, the price for bidder 1 to buy g_1 is $\max(v^{-1}_{\{g_1\}}, v^{-1}_{\{g_1,\ldots,g_m\}}) = \max(0,5) = 5$. Thus, he buys g_1 at 5. On the other hand, the price for bidder 2 to buy g_2 is $\max(v^{-2}_{\{g_1\}}, v^{-2}_{\{g_1,\ldots,g_m\}}) = \max(6,5) = 6$. Thus, he buys nothing. In this case, only g_1 is allocated. In VCG, if bidder 1 and 2 are actually false-name bids of a single bidder whose valuation for both goods is 7, the bidder can decrease his payment from 5 to $3+0=3$. Meanwhile, in ARP, submitting false-name bids does not decrease his payment *and* he wins only g_1.

We now show the key properties of ARP for single-minded bidders.

Theorem 5 *ARP with m different goods satisfies allocation feasibility for single-minded bidders.*

PROOF. Let us assume bidders $0,1,\ldots,m$ submit $v^*_{\{g_1,\ldots,g_m\}}$, $v^*_{\{g_1\}}$, $\ldots$, $v^*_{\{g_m\}}$, respectively. For bidders other than these, it is clear that they cannot obtain any positive utility. Thus, to show allocation feasibility, it suffices to show that if bidder 0 wants $\{g_1,\ldots,g_m\}$, then no other bidder $1,\ldots,m$ can obtain positive utilities.

Assume that bidder 0's utility is positive when he obtains $\{g_1,\ldots,g_m\}$. For each $1 \leq k \leq m$, $v^{-0}_{g_{\{(k)\}}} = v^*_{\{g_k\}}$ holds. From the definition of ARP, we obtain

$$\begin{aligned} v^*_{\{g_1,\ldots,g_m\}} &> p(\{g_1,\ldots,g_m\},\Theta_{-0}), \\ &= \max(v^{-0}_{\{g_1,\ldots,g_m\}}, v^{-0}_{\{g_{(1)}\}}, 2v^{-0}_{\{g_{(2)}\}}) \\ &= \max(v^{-0}_{\{g_1,\ldots,g_m\}}, v^*_{\{g_1\}}, 2v^*_{\{g_2\}}). \end{aligned}$$

Thus, $v^*_{\{g_1,\ldots,g_m\}} > v^*_{\{g_1\}}$ and $v^*_{\{g_1,\ldots,g_m\}} > 2v^*_{\{g_2\}}$ hold.

We first show that for bidder 1, who is bidding $v^*_{\{g_1\}}$, his utility when obtaining $\{g_1\}$ is negative. Thus, $p(\{g_1\}, \Theta_{-1}) > v^*_{\{g_1\}}$ holds. Since we assume bidders are single-mined, $v^{-1}_{\{g_1,\ldots,g_m\}} = v^*_{\{g_1,\ldots,g_m\}}$ holds. When $v^{-1}_{\{g_1\}} \geq v^*_{\{g_2\}}$ holds, $v^{-1}_{\{g_{(1)}\}} = v^{-1}_{\{g_1\}}$ holds. Also, for each $2 \leq k \leq m$, $v^{-1}_{\{g_{(k)}\}} = v^*_{\{g_k\}}$ holds. Since $v^*_{\{g_1,\ldots,g_m\}} > 2v^*_{\{g_2\}}$ holds, we obtain

$$\begin{aligned} p(\{g_1\}, \Theta_{-1}) = & \max(v^{-1}_{\{g_{(1)}\}}, v^{-1}_{\{g_1,\ldots,g_m\}}) \\ \geq & \ v^*_{\{g_1,\ldots,g_m\}} > v^*_{\{g_1\}}. \end{aligned}$$

On the other hand, when $v^{-1}_{\{g_1\}} < v^*_{\{g_2\}}$ holds, $v^{-1}_{\{g_{(1)}\}} = v^*_{\{g_2\}}$ also holds. Since $v^*_{\{g_1,\ldots,g_m\}} > 2v^*_{\{g_2\}}$ holds, we obtain

$$\begin{aligned} p(\{g_1\}, \Theta_{-1}) = & \max(v^{-1}_{\{g_{(1)}\}}, v^{-1}_{\{g_1,\ldots,g_m\}}) \\ \geq & \ v^*_{\{g_1,\ldots,g_m\}} > v^*_{\{g_1\}}. \end{aligned}$$

Accordingly, bidder 1 cannot obtain positive utility by obtaining $\{g_1\}$.

Next, for any bidder k who bids $v^*_{\{g_k\}}$, where $2 \leq k \leq m$, we show that his utility when obtaining $\{g_k\}$ is non-positive. Thus, $p(\{g_k\}, \Theta_{-k}) \geq v^*_{\{g_k\}}$ holds. From our assumptions, $v^{-k}_{\{g_{(1)}\}} = v^*_{\{g_1\}}$ holds. When $v^*_{\{g_1,\ldots,g_m\}}/2 < v^*_{\{g_1\}}$,

$$p(\{g_k\}, \Theta_{-k}) \geq v^*_{\{g_1,\ldots,g_m\}}/2 > v^*_{\{g_2\}} \geq v^*_{\{g_k\}}$$

holds. On the other hand, when $v^*_{\{g_1,\ldots,g_m\}}/2 \geq v^*_{\{g_1\}}$,

$$p(\{g_k\}, \Theta_{-k}) \geq v^*_{\{g_1,\ldots,g_m\}} > 2v^*_{\{g_2\}} > v^*_{\{g_k\}}$$

hold. Accordingly, bidder k cannot obtain positive utility by obtaining $\{g_k\}$. □

As shown in [15], a PORF mechanism that satisfies allocation feasibility is automatically strategy-proof. Thus, ARP is strategy-proof for single-minded bidders.

Theorem 6 *ARP with m different goods is false-name-proof for single-minded bidders.*

PROOF. Let us assume bidders $0, 1, \ldots, m$ submit $v^*_{\{g_1,\ldots,g_m\}}$, $v^*_{\{g_1\}}$, ..., $v^*_{\{g_m\}}$, respectively. For bidders other than 0, 1, and 2, it is clear that they cannot achieve positive utilities even if they use false-name bids, since their prices never fall below $v^*_{g_2}$. Thus, it suffices to consider the possibilities of false-name bids by bidder $0, 1$, and 2 only.

First, let us consider the possibility of a false-name bid by bidder 0 and assume bidder 0 is actually interested in bundle $\{g_i, g_j\}$ for i and $j (> i)$. Assume he uses identifiers i' and j', and bids $v^{*'}_{\{g_i\}}$ and $v^{*'}_{\{g_j\}}$, respectively, so that they become winners. W.l.o.g., we assume $v^{*'}_{\{g_i\}} \geq v^{*'}_{\{g_j\}}$ holds. To be winners, in addition to $v^{-0}_{\{g_1,\ldots,g_m\}} < 2v^*_{\{g_j\}}$, $v^{*'}_{\{g_i\}} \geq v^*_{\{g_1\}}$ and $v^{*'}_{\{g_j\}} \geq v^*_{\{g_2\}}$ must hold. Thus, it is clear that $p(\{g_i\}, \Theta_{-0} \cup \{j'\}) \geq v^*_{\{g_1\}}$ and $p(\{g_j\}, \Theta_{-0} \cup \{i'\}) \geq v^*_{\{g_2\}}$ hold. The sum of the prices given to i' and j' is

$$\begin{aligned} & p(\{g_i\}, \Theta_{-0} \cup \{j'\}) + p(\{g_j\}, \Theta_{-0} \cup \{i'\}) \\ \geq & \max(v^*_{\{g_1\}}, v^{-0}_{\{g_1,\ldots,g_m\}}/2) + \max(v^*_{\{g_2\}}, v^{-0}_{\{g_1,\ldots,g_m\}}/2) \\ \geq & \max(v^*_{\{g_1\}}, v^{-0}_{\{g_1,\ldots,g_m\}}) = p(\Theta_{-0}, \{g_1, \ldots, g_m\}). \end{aligned}$$

Thus, bidder 0 cannot decrease the payment using false-name bids.

Next, we consider the possibility of false-name bids by bidder 1. When bidder 1 is winning and $v^*_{\{g_1,\ldots,g_m\}} < 2v^*_{g_2}$ holds, his price for $\{g_1\}$ is given by

$$p(\{g_1\}, \Theta_{-1}) = \max(v^{-1}_{\{g_1\}}, v^*_{\{g_2\}}, v^*_{\{g_1,\ldots,g_m\}}/2).$$

This price cannot be decreased by adding false-name bids. On the other hand, when $v^*_{\{g_1,\ldots,g_m\}} \geq 2v^*_{g_2}$ holds,

$$p(\{g_1\}, \Theta_{-1}) = \max(v^{-1}_{\{g_1\}}, v^*_{\{g_2\}}, v^*_{\{g_1,\ldots,g_m\}}).$$

This price can be manipulated if he submits a false-name bid $v^{*'}_{g_2}$ on $\{g_2\}$, using another identifier 2', so that $v^{*'}_{g_2} > v^*_{\{g_1,\ldots,g_m\}}/2$ holds. Then he can decrease his price; we obtain

$$p(\{g_1\}, \Theta_{-1} \cup \{2'\}) = \max(v^{-1}_{\{g_{(1)}\}}, v^{-1}_{\{g_1,\ldots,g_m\}}/2).$$

Since bidder 1 already obtains $\{g_1\}$ when not using a false-name bid, $v^*_{\{g_1\}} > v^*_{\{g_1,\ldots,g_m\}}$ must hold. Then, for 2',

$$p(\{g_2\}, \Theta_{-2'}) = \max(v^*_{g_2}, v^*_{\{g_1,\ldots,g_m\}}/2).$$

Clearly, $v^{*'}_{g_2} > v^*_{g_2}$ holds. Since, from the assumptions, $v^{*'}_{g_2} > v^*_{\{g_1,\ldots,g_m\}}/2$ holds, we obtain $p(\{g_2\}, \Theta_{-2'}) < v^{*'}_{g_2}$. As a result, bidder 2' must be a winner, and the total payment is larger than, or at least equal to, the payment when not using a false-name bid. If bidder 1 uses one more identifier $3'$ and submits $v^{*'}_{\{g_3\}}$, which is equal to $v^{*'}_{\{g_2\}}$ to $\{g_3\}$, then

$$p(\{g_2\}, \Theta_{-2'}) = p(\{g_3\}, \Theta_{-3'}) = v^{*'}_{\{g_2\}} = v^{*'}_{\{g_3\}}.$$

In this case, the utilities of bidder $2'$ and $3'$ are 0. However, since ARP chooses an allocation that has two winners whenever possible, either bidder $2'$ or $3'$ becomes a winner. Thus, for bidder 1, it is impossible to decrease his payment by submitting false-name bids.

Finally, we consider the possibility of false-name bids by bidder 2. If bidder 2 is winning, his payment is

$$p(\{g_2\}, \Theta_{-2}) = \max(v^{-2}_{\{g_2\}}, v^*_{\{g_3\}}, v^*_{\{g_1,\ldots,g_m\}}/2).$$

This price cannot be decreased by adding false-name bids. □

Theorem 7 *For ARP with m different goods, the worst-case efficiency ratio is $\frac{2}{m+1}$ for single-minded bidders.*

PROOF. ARP allocates m goods according to the declared bids in the three following ways: Case i: $\{g_1, \ldots, g_m\}$ is allocated to the bidder who bids $v^*_{\{g_1,\ldots,g_m\}}$, Case ii: Only good g_1 is allocated to the bidder who bids $v^*_{\{g_1\}}$, Case iii: good g_1 is allocated to the bidder who bids $v^*_{\{g_1\}}$ and good g_2 is allocated to the bidder who bids $v^*_{\{g_2\}}$. Let us examine the worst-case efficiency ratio in those cases.

In Case i, $v^*_{\{g_1,\ldots,g_m\}} \geq v^*_{\{g_1\}}$ and $v^*_{\{g_1,\ldots,g_m\}} \geq 2v^*_{\{g_2\}}$ hold, and the social surplus s^{ARP} that ARP achieves is $v^*_{\{g_1,\ldots,g_m\}}$. In this case, let us show an upper bound on the social surplus s^* that the Pareto efficient allocation achieves.

Let us examine each bid v_B in the Pareto efficient allocation. First, let us consider the case where $|B| \geq 2$. Then, $v_B \leq v^*_{\{g_1,\ldots,g_m\}}$ holds (otherwise, v_B must be the largest bid for all goods). Thus, $v_B/|B| \leq v^*_{\{g_1,\ldots,g_m\}}/2$ holds, i.e., the average value per good is at most $v^*_{\{g_1,\ldots,g_m\}}/2$. Next, let us consider the case where $|B| = 1$. If $B = \{g_1\}$, $v_B \leq v^*_{\{g_1\}} \leq v^*_{\{g_1,\ldots,g_m\}}$ holds. Otherwise, $v_B \leq v^*_{\{g_2\}} \leq v^*_{\{g_1,\ldots,g_m\}}/2$ holds. Since at most m goods belong to the Pareto efficient allocation, $s^* \leq v^*_{\{g_1,\ldots,g_m\}} + (m-1)v^*_{\{g_1,\ldots,g_m\}}/2$ holds. Thus, we obtain

$$\frac{s^{ARP}}{s^*} \geq \frac{v^*_{\{g_1,\ldots,g_m\}}}{v^*_{\{g_1,\ldots,g_m\}}+(m-1)v^*_{\{g_1,\ldots,g_m\}}/2} = \frac{2}{m+1}.$$

In Case ii, $v^*_{\{g_1,\ldots,g_m\}} < v^*_{\{g_1\}}$ and $v^*_{\{g_1,\ldots,g_m\}} \geq 2v^*_{\{g_2\}}$ hold and s^{ARP} is $v^*_{\{g_1\}}$. Let us examine each bid v_B in the Pareto efficient allocation. First, let us consider the case where $|B| \geq 2$. Then, $v_B \leq v^*_{\{g_1\}}$ holds (otherwise, v_B must be the largest bid for all goods and $v^*_{\{g_1,\ldots,g_m\}} < v^*_{\{g_1\}}$ does not hold). Thus, $v_B/|B| \leq v^*_{\{g_1\}}/2$ holds, i.e., the average value per good is at most $v^*_{\{g_1\}}/2$.

Next, let us consider the case where $|B| = 1$. If $B = \{g_1\}$, $v_B \leq v^*_{\{g_1\}}$ holds. Otherwise, $v_B \leq v^*_{\{g_2\}}$ holds. In this case, from $v^*_{\{g_1,\ldots,g_m\}} < v^*_{\{g_1\}}$ and $v^*_{\{g_1,\ldots,g_m\}} \geq 2v^*_{\{g_2\}}$, $v_B \leq v^*_{\{g_2\}} < v^*_{\{g_1\}}/2$ holds. Since at most m goods belong to the Pareto efficient allocation, $s^* \leq v^*_{\{g_1\}}+(m-1)v^*_{\{g_1\}}/2$ holds. Thus, we obtain

$$\frac{s^{ARP}}{s^*} \geq \frac{v^*_{\{g_1\}}}{v^*_{\{g_1\}}+(m-1)v^*_{\{g_1\}}/2} \geq \frac{2}{m+1}.$$

In Case iii, $v^*_{\{g_1,\ldots,g_m\}} < 2v^*_{\{g_2\}}$ holds, and s^{ARP} is $v^*_{\{g_1\}} + v^*_{\{g_2\}}$. Let us examine each bid v_B in the Pareto efficient allocation. First, let us consider the case where $|B| \geq 2$. Then, $v_B \leq v^*_{\{g_1,\ldots,g_m\}}$ holds (otherwise, v_B must be the largest bid for all goods). Since $v^*_{\{g_1,\ldots,g_m\}} < 2v^*_{\{g_2\}}$ holds, $v_B/|B| < v^*_{\{g_2\}}$ holds, i.e., the average value per good is at most $v^*_{\{g_2\}}$. Next, let us consider the case where $|B| = 1$. If $B = \{g_1\}$, $v_B \leq v^*_{\{g_1\}}$ holds. Otherwise, $v_B \leq v^*_{\{g_2\}}$ holds. Since at most m goods belong to s^*, $s^* \leq v^*_{\{g_1\}}+(m-1)v^*_{\{g_2\}}$ holds. Thus, we obtain

$$\begin{aligned}\frac{s^{ARP}}{s^*} &\geq \frac{v^*_{\{g_1\}}+v^*_{\{g_2\}}}{v^*_{\{g_1\}}+(m-1)v^*_{\{g_2\}}} \\ &= \frac{2+(v^*_{\{g_1\}}-v^*_{\{g_2\}})/v^*_{\{g_2\}}}{m+(v^*_{\{g_1\}}-v^*_{\{g_2\}})/v^*_{\{g_2\}}} \geq \frac{2}{m},\end{aligned}$$

which is greater than $\frac{2}{m+1}$.

Therefore, the worst-case efficiency ratio that ARP achieves is at least $\frac{2}{m+1}$. Also, from Theorem 1, it cannot be more than $\frac{2}{m+1}$. More specifically, let us assume the situation in Case 9, where $m = k+1$. In a Pareto efficient mechanism, each $g_k \in [1,m]$ is allocated to bidder k. The obtained social surplus is $(c/2-\epsilon)m+c/2$. On the other hand, in ARP, both goods are allocated to bidder 0. The obtained social surplus is c. Then, the ratio is $\frac{2}{(1+\epsilon)m+1}$, which can be made arbitrarily close to $2/(m+1)$. Thus, the worst-case efficiency ratio that ARP achieves is $\frac{2}{m+1}$. □

The ARP mechanism is somewhat limited, i.e., it allocates goods to at most two bidders. However, extending ARP so that it can allocate three or more individual goods is not easy. It might be impossible considering that the current ARP already achieves the optimal worst-case efficiency ratio.

6. CONCLUSION

We showed that the worst-case efficiency ratio of any false-name-proof mechanism that satisfies some apparently minor assumptions is at most $\frac{2}{m+1}$. We also showed that the worst-case efficiency ratio of existing false-name-proof mechanisms is $\frac{1}{m}$ or 0. Furthermore, we developed an optimal mechanism, ARP, which is false-name-proof when all bidders are single-minded. When bidders are not necessarily single-minded, there still exists a gap between the upper bound ($\frac{2}{m+1}$) and the best ratio obtained by existing mechanisms ($\frac{1}{m}$). In our future work, we hope to eliminate this gap, either by developing a mechanism with a better ratio, or by proving a tighter upper bound.

Acknowledgement

Iwasaki and Yokoo are partially supported by the Ministry of Education, Science, Sports and Culture, Grant-in-Aid for Scientific Research (A), 20240015, 2008. Conitzer and Guo are supported by NSF under award number IIS-0812113, and by an Alfred P. Sloan Research Fellowship.

7. REFERENCES

[1] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *FOCS*, pages 295–304, 2004.

[2] A. Archer and Éva Tardos. Frugal path mechanisms. In *SODA*, pages 991–999, 2002.

[3] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2005.

[4] A. Czumaj and A. Ronen. On the expected payment of mechanisms for task allocation. In *PODC*, pages 98–106, 2004.

[5] S. Dobzinski and N. Nisan. Mechanisms for multi-unit auctions. In *EC*, pages 346–351, 2007.

[6] E. Elkind, A. Sahai, and K. Steiglitz. Frugality in path auctions. In *SODA*, pages 701–709, 2004.

[7] M. Guo and V. Conitzer. Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69–98, 2009.

[8] A. R. Karlin, D. Kempe, and T. Tamir. Beyond VCG: Frugality of truthful mechanisms. In *FOCS*, pages 615–624, 2005.

[9] R. Lavi, A. Mu'alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *FOCS*, pages 574–583, 2003.

[10] J. O. Ledyard. Optimal combinatoric auctions with single-minded bidders. In *EC*, pages 237–242, 2007.

[11] H. Moulin. Efficient, strategy-proof and almost budget-balanced assignment. *Journal of Economic Theory*, forthcoming, 2009.

[12] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[13] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[14] T. Roughgarden and E. Tardos. How bad is selfish routing? In *FOCS*, pages 236–259, 2000.

[15] M. Yokoo. The characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. In *IJCAI*, pages 733–739, 2003.

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

[17] 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.