

Multi-Unit Auction Revenue with Possibilistic Beliefs

by

Georgios Vlachos

Submitted to the Department of Electrical Engineering and Computer Science

in partial fulfillment of the requirements for the degree of

Master of Engineering in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2017

© Massachusetts Institute of Technology 2017. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
August 15, 2017

Certified by
Silvio Micali
Ford Professor of Engineering
Thesis Supervisor

Accepted by
Christopher J. Terman
Chairman, Department Committee on Graduate Theses

Multi-Unit Auction Revenue with Possibilistic Beliefs

by

Georgios Vlachos

Submitted to the Department of Electrical Engineering and Computer Science
on August 15, 2017, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Computer Science and Engineering

Abstract

The revenue of traditional auction mechanisms is benchmarked solely against the players' own valuations, despite the fact that they may also have valuable beliefs about each other's valuations.

Not much is known about generating revenue in auctions of multiple identical copies of a same good. (In particular the celebrated Vickrey mechanism has no revenue guarantees.)

For such auctions, we (1) put forward an attractive revenue benchmark, based on the players' possibilistic about each other, and (2) construct a mechanism that achieves such benchmark, assuming that the players are two-level rational (where the rationality is in the sense of Aumann).

Thesis Supervisor: Silvio Micali

Title: Ford Professor of Engineering

1 Introduction

Generating revenue has always been a main desideratum in auctions. Guaranteeing higher revenue is possible when the seller has significant information about the players. For instance, in auctions of a single good, Myerson shows how to construct optimal auctions when the seller knows, for each player i , the distribution D_i from which i 's valuation has been drawn.

In single-good auctions, however, a significant revenue benchmark can be guaranteed even when the seller has no information about the valuations of the players. Indeed, the second-price mechanism guarantees that in (weakly) dominant strategies the revenue generated coincides with the second highest valuation for the good for sale.

The Problem of Revenue in Multi-Unit Auctions In multi-unit auctions, m identical copies of the same good are available for sale. Not much is known about revenue in such auctions.

In particular, the celebrated Vickrey mechanism [13] yields maximum social welfare in (weakly) dominant strategies, but offers no guarantees about revenue. For instance, if all players value little a second copy of the good and at most m players significantly value a first copy of the good, then the Vickrey mechanism will not generate any significant revenue.

Our Goal We wish to investigate generating revenue, in multi-unit auctions, when the seller has no information about the players, but the players have significant information about each other's valuations. Since we allow the players to have wrong information about the valuations of their opponents, we should talk about *beliefs*.

Possibilistic Beliefs In our auctions, a valuation v consists of an m -tuple of non-negative reals, $(v(1), \dots, v(m))$, where $v(k)$ represents the marginal value of receiving a k th copy of the good. Following Vickrey, we assume decreasing marginal valuations: that is, $v(1) \geq v(2) \geq \dots \geq v(m)$.

As usual, a player i knows his own true valuation, θ_i^* . In addition, he may also have some belief about the true valuations of his opponents, $-i$. For instance, i 's belief may consist of a probability distribution from which i believes that the true valuation subprofile of his opponents, θ_{-i}^* , has been drawn. However, i 's belief may be less structured. In particular, i may be unable to determine how likely it is for a specific subprofile θ_{-i} to be the true one. Furthermore, he may even be unable to estimate the relative likelihood of two subprofiles θ_{-i} and θ'_{-i} . Accordingly, following [6, 7, ?], we consider quite elementary beliefs.

Essentially, a player i 's belief is just a *set* consisting of all the valuation subprofiles θ_{-i} that i considers possible for his opponents. As we shall prove, it is possible to use even such elementary beliefs to generate revenue.

Revenue Benchmark Our revenue target is the *collective external-belief* (CEB, for short). This benchmark is defined in our technical sections. Intuitively, CEB

represents the highest social welfare one can guarantee by pooling together the beliefs of all players about their opponents.

The CEB benchmark is more easily described for the unit-demand setting, that is, when each player i may positively value a first copy of the good, but has zero value for any extra copies: $\theta_i^*(1) \geq 0$ while $\theta_i^*(2) = \dots = \theta_i^*(m) = 0$.¹ In this case, CEB is at least $v_1^* + \dots + v_m^*$ if there exists an allocation of the m copies in which, for every player i receiving a copy, at least another player j believes that i 's value for the copy received is at least v_i^* . Notice, therefore, that the benchmark CEB is very demanding, as it “pools together” the beliefs of all the players about their opponents.

Solution Concept We assume that each player is rational in a minimal sense: namely, that, based on his beliefs, he never plays a strictly dominated action, and believes that his opponents behave in the same way. In other words, we assume that the players are two-level rational, where the notion of rationality is in the sense of Aumann [1].

Main Results We put forward a normal-form mechanism that, when played by such two-level rational players, yields a revenue that, as the number of copies increases, approaches the CEB benchmark.

We first exhibit, in Section 4, a mechanism for the unit-demand setting. This mechanism virtually achieves a fraction $\left(1 - \frac{\lceil \sqrt{2m} \rceil}{m}\right)$ of CEB.

Then, in Section 5, we generalize this first mechanism to the q -unit demand setting. The generalized mechanism virtually achieves a fraction $\left(1 - \frac{q(\lceil \sqrt{\frac{2m}{q}} \rceil + 1)}{m}\right)$ of CEB.

1.1 Prior Work On Possibilistic Beliefs

The traditional probabilistic approach to modeling the players' beliefs in settings of incomplete information was pioneered by Harsanyi [10]. An informative survey about this approach has been given by Siniscalchi [12].

Following Aumann [1], we instead take a set-theoretic approach to model the players' beliefs. Actually, Kripke independently studied set-theoretic models of beliefs using modal logic [11]. (For an alternative description of Kripke's approach see [8].)

More recently, possibilistic beliefs have been used to generate revenue in auctions different from ours. First-order possibilistic beliefs have been used by Chen and Micali in single-good auctions [6]. Also in single-good auctions, higher-order possibilistic beliefs have been used by Chen, Micali, and Pass [7]. The authors of another paper, reference omitted, use possibilistic beliefs, so as to generate revenue in combinatorial auctions under essentially the solution concept adopted in this paper.² Since multi-unit auctions can be considered as a special case of combinatorial ones, let us emphasize that their mechanism, run in multi-unit auctions, achieves (within a factor

¹For instance, if the m goods sold in the auctions are non-exclusive licenses under a given patent, then the setting is naturally unit-demand, as no one needs two non-exclusive licence under the same patents.

²They assume that the players will only consider strategies surviving two levels of elimination of weakly dominated strategies (based on their beliefs). We instead assume that the players stick to strategies surviving two levels of elimination of strictly dominated strategies.

of 2) a revenue benchmark, BB , less demanding than CEB . Essentially, they define BB_i , the best belief of a player i , to be the maximum social welfare i may guarantee, based on his own beliefs about his opponents, by allocating the copies of the goods to just his opponents. Then, they define BB to be the maximum of the BB_i 's. By contrast, CEB is the maximum social welfare one could guarantee, based only on all the players' beliefs about their opponents, by allocating the copies of the good to the set of all players.

Finally, let us mention the work of Bergemann and Morris on robust mechanism design [3]. Their work is related to ours in the sense that they also consider implementing social-choice correspondences defined on more than just the players' payoff types. However, their goals are different from ours. While we design alternative mechanisms that use the players' beliefs about their opponents, they use these beliefs to strengthen the meaningfulness of implementation in dominant strategies. More recently, Bergemann and Morris [4] have pointed out that social choice correspondences defined solely over the player's payoff types cannot represent revenue-maximizing allocations.

2 Basic Notions

2.1 Multi-Unit Auctions

In this paper we only consider *multi-unit* auctions. In such auctions there are finitely many identical copies of the same good for sale. We shall denote the actual number of such copies by m , and let $[m] = \{1, \dots, m\}$. The set of all valuations is Θ . A valuation v is a function mapping 0 to 0, and every integer in $\{1, \dots, m\}$ to a non-negative real number, representing the marginal value for a j th copy. (I.e., $v(j)$ is the value associated to the j th copy to one who already has $j - 1$ copies.) Following Vickery [?], we assume "decreasing-marginal valuations", that is, $v(1) \geq \dots \geq v(m) \geq 0$.

The set of players is $N = \{1, 2, \dots, n\}$, and the true valuation of player i is θ_i^* . An outcome ω is a pair (A, P) , where $A = (A_0, \dots, A_n)$ is a vector of nonnegative integers and P is a profile of real numbers. Vector A is referred to as the allocation of ω and must be such that $\sum_{i=0}^n A_i = m$. Component A_0 represents the number of unallocated copies and, for each $i > 0$, A_i represents the number of copies allocated to player i . Each P_i represents the price paid by player i .

For each player i , i 's utility function u_i maps a valuation θ_i and an outcome $\omega = (A, P)$ to $u_i(\theta_i, \omega) \triangleq \sum_{j=1}^{A_i} \theta_i(j) - P_i$. The true utility of player i for an outcome ω is $u_i(\theta_i^*, \omega)$. When θ_i^* is clear, we may write $u_i(\omega)$ instead of $u_i(\theta_i^*, \omega)$.

The *individual welfare* of player i in an allocation A is $\sum_{j=1}^{A_i} \theta_i^*(j)$. The set of all possible allocations is denoted by \mathcal{A} . The *maximum social welfare* is $\max_{A \in \mathcal{A}} \sum_i \sum_{j=1}^{A_i} \theta_i^*(j)$.

A *mechanism* M specifies:

- For each player i , the set S_i of pure strategies available to i .
(The set of all strategy profiles is $S = S_1 \times \dots \times S_n$.)
- A function mapping each strategy profile $s \in S$ to an outcome, $M(s)$.
(Accordingly, $u_i(M(s))$ is i 's true utility for the outcome $M(s)$.)

If M is probabilistic, $M(s)$ is a distribution over outcomes, and $u_i(M(s))$ is the corresponding expected utility of i . If the underlying mechanism M is clear, we may write $u_i(s)$ instead of $u_i(M(s))$.

A mechanism M is *individually rational* (IR), if every player i has a strategy $safe_i \in S_i$ such that for all strategy profiles s for which $s_i = safe_i$, $u_i(M(s)) \geq 0$.

2.2 Beliefs

We model a player i 's beliefs as a set, \mathbb{B}_i , which in i 's view, consists of *all possible candidates for the true valuation profile*. Since i knows his own true valuation, θ_i^* , the i th component of each valuation profile in \mathbb{B}_i must coincide with θ_i^* . That is,

Definition 1. For each player i , $\mathbb{B}_i \subset \{\theta \in \Theta : \theta_i = \theta_i^*\}$. $\mathbb{B} = \mathbb{B}_1 \times \dots \times \mathbb{B}_n$.

Note that the true valuation profile θ^* need not belong to \mathbb{B}_i , in which case we may say that i 's beliefs are wrong.

Also note that \mathbb{B}_i need not capture all beliefs of i . For instance, i may have the following vague probabilistic belief about some opponent j : j 's true valuation is θ_j with some probability between 0 and 1/2. However, all “additional beliefs of i must not contradict \mathbb{B}_i ”. For example, if i had the additional belief just mentioned, then \mathbb{B}_i must contain at least one valuation profile t such that $t_j = \theta_j$.

2.3 The CEB Benchmark

The collective external belief benchmark is more simply stated for the unit-demand case.

Definition 2. Let \mathbb{B} be a belief profile for a multi-unit auction with unit-demand; $v_i^{\mathbb{B}} \triangleq \max_{j \neq i} \min_{\theta \in \mathbb{B}_j} \theta_i$; and $P = (p_1, p_2, \dots, p_n)$ a permutation of the players such that $v_{p_1}^{\mathbb{B}} \geq v_{p_2}^{\mathbb{B}} \geq \dots$. Then,

$$CEB = \sum_{i=1}^m v_{p_i}^{\mathbb{B}} .$$

Notice that $v_i^{\mathbb{B}}$ is the highest value v such that at least one opponent of i is sure (according to his belief \mathbb{B}_i) that i values a copy of the good least v . Thus, CEB essentially aggregates, in the best possible way, the knowledge that the players have about their opponents.

Let us now generalize CEB to the q -unit demand case, that is when the marginal value of each player for a $(q + 1)$ st copy of the good is 0.

Definition 3. Let \mathbb{B} be a belief profile for a multi-unit auction and \mathcal{A} the set of all possible allocations. For every $\text{copy} \in [m]$ define

$$v_i^{\mathbb{B}}[\text{copy}] \triangleq \max_{j \neq i} \min_{\theta \in \mathbb{B}_j} \theta_i(\text{copy}) .$$

Then,

$$CEB = \max_{A \in \mathcal{A}} \sum_{i=1}^n \sum_{\text{copy}=1}^{A_i} v_i^{\mathbb{B}}[\text{copy}] .$$

Notice that $v_i^{\mathbb{B}}[j]$ is the highest value v such that at least an opponent j of i is sure (according to his beliefs \mathbb{B}_j) that i values a j th copy of the good at least v .

Also notice that CEB is the maximum social welfare if the true valuation of each player i were the function mapping each integer $j \in [m]$ to $v_i^{\mathbb{B}}(j)$.

On the Usefulness of Possibilistic Beliefs Traditionally, revenue benchmarks (and more generally social choice functions) are defined relative to the true valuation profile. Quite differently, the collective external belief benchmark is defined relative to the players' beliefs, more precisely relative to their external beliefs. The ability of defining revenue benchmarks (and more generally social choice functions) over the players' beliefs enlarges the set of meaningful "targets" available to a mechanism designer.

On the Significance of CEB At times, the mechanism designer may have little or no information about the players, while the players may know each other quite well. When the latter is the case, the CEB benchmark is high. For instance, if the players knew each other's valuations at least —say— 80% correctly, then the CEB benchmark would be at least 80% of the maximum social welfare.

By saying that "the players know each other valuations at least 80% correctly" we mean that, for all players i and copies c , there is a value $v_{i,c}$ and an opponent j of i such that j is sure that $v_{i,c} \geq \theta_i^*(c)$.³

When some player has unrealistically high beliefs about the valuations of one or more of his opponents, it is possible that CEB is much higher than the maximum social welfare.

Examples Let us illustrate the benchmark CEB by means of the following two examples. In each of them, there are 3 players and 2 copies of the good.

EXAMPLE 1. In this example, each player happens to believe that only three valuation profiles are possible.

$$\mathbb{B}_1 = \{ ((100, 80), (85, 83), (90, 68)), ((100, 80), (85, 80), (89, 63)), ((100, 80), (88, 84), (92, 71)) \}$$

$$\mathbb{B}_2 = \{ ((94, 76), (88, 84), (90, 60)), ((102, 84), (88, 84), (86, 58)), ((100, 80), (88, 84), (90, 64)) \}$$

$$\mathbb{B}_3 = \{ ((105, 90), (90, 80), (90, 68)), ((100, 84), (88, 79), (90, 68)), ((98, 77), (86, 78), (90, 68)) \}$$

Player 1's set of beliefs implies that $\theta_2(1)^* \geq 85$, $\theta_2(2)^* \geq 80$, $\theta_3(1)^* \geq 89$, and $\theta_3(2)^* \geq 63$.

Player 2's set of beliefs implies that $\theta_1(1)^* \geq 94$, $\theta_1(2)^* \geq 76$, $\theta_3(1)^* \geq 86$, and $\theta_3(2)^* \geq 58$.

Player 3's set of beliefs implies that $\theta_1(1)^* \geq 98$, $\theta_1(2)^* \geq 77$, $\theta_2(1)^* \geq 86$, and $\theta_2(2)^* \geq 78$.

Thus in example 1, $CEB = 98 + 89 = 187$. (This value of CEB is actually realized because player 3 believes that player 1 values a copy at ≥ 98 and player 1 believes that player 3 values a copy at ≥ 89 .)

³Actually, there is no need that the players have at least 80% correct knowledge about each possible player and each possible copy. It is enough that they have such knowledge about "the players and copies that count".

For comparison, notice that in this example, the maximum social welfare is $100 + 90 = 190$. (In fact, $\theta_1^* = (100, 80)$, $\theta_2^* = (88, 84)$, and $\theta_3^* = (90, 68)$, and the maximum social welfare is achieved by allocating one copy to player 1, and another copy to player 3.)

Moreover, the Vickrey mechanism generates revenue $88 + 88 = 176$. \triangle

EXAMPLE 2:

$$\mathbb{B}_1 = (10, 8) \times ([12, 14] \times [6, 8]) \times ([12, 15] \times [10, 11])$$

$$\mathbb{B}_2 = ([9, 11] \times [7, 9]) \times (10, 7) \times ([8, 12] \times [6, 7])$$

$$\mathbb{B}_3 = ([10, 13] \times [9, 10]) \times ([8, 12] \times [6, 8]) \times (7, 5)$$

Player 1's set of beliefs implies that $\theta_2(1)^* \geq 12$, $\theta_2(2)^* \geq 6$, $\theta_3(1)^* \geq 12$, and $\theta_3(2)^* \geq 10$.

Player 2's set of beliefs implies that $\theta_1(1)^* \geq 9$, $\theta_1(2)^* \geq 7$, $\theta_3(1)^* \geq 8$, $\theta_3(2)^* \geq 6$.

Player 3's set of beliefs implies that $\theta_1(1)^* \geq 10$, $\theta_1(2)^* \geq 9$, $\theta_2(1)^* \geq 8$, $\theta_2(2)^* \geq 6$.

In example 2, $CEB = 12 + 12 = 24$. (This value of CEB is actually realized because player 1 believes that player 2 values a first copy at ≥ 12 and he also believes that player 3 values a first copy at ≥ 12 .)

The maximum social welfare is $10 + 10$. (In fact, $\theta_1^* = (10, 8)$, $\theta_2^* = (10, 7)$, and $\theta_3^* = (7, 5)$, and the maximum social welfare is achieved by allocating one copy to player 1, and another copy to player 2.)

Moreover, the Vickrey mechanism generates revenue $7 + 8 = 15$.

Example 2 shows that the CEB benchmark can exceed the maximum social welfare when the players have incorrect beliefs about their opponents' valuations. In our case, player 1 believes that player 2 values a first copy of the good at least 12 and that player 3 values a first copy of the good at least 12. However, player 2 values a first copy of the good at 10 and player 3 at 7. \triangle

3 Solution Concept

We consider implementation in *strictly conservative strategies*, as defined [6]. Essentially, we rely on the players to perform two levels of deletion of strictly dominated strategies, "based on their beliefs".

Recall that iterated elimination of strictly dominated strategies is traditionally defined in setting of complete information, where each player i knows the true valuations (more generally payoff type) of his opponents. It is thus necessary to generalize this notion to our incomplete-information setting.

Recall that the set of probability distributions over a set X is denoted by $\Delta(X)$.

Definition 4 (Dominated and Undominated Strategies). *Let i be a player, θ_i a valuation of i , and $T = T_i \times T_{-i}$ a subset of strategy profiles. A strategy $s_i \in T_i$ is strictly dominated by another strategy $s'_i \in T_i$, relative to θ_i and T , in symbols $s_i <_T^{\theta_i} \sigma_i$, if, for all strategy subprofiles $s_{-i} \in T_{-i}$,*

$$u_i(\theta_i, (s_i, s_{-i})) < u_i(\theta_i, (s'_i, s_{-i})).$$

We let $U_i(\theta_i, T)$ denote the set of strategies in T_i that are not strictly dominated relative to θ_i and T .

Accordingly, knowing that his true valuation is θ_i^* and that the initial set of strategy profiles is S , a rational player i eliminates some of his initial strategies and will only consider strategies in

$$U_i \triangleq U_i(\theta_i^*, S).$$

Based on his belief \mathbb{B}_i and believing that his opponents are rational, i will eliminate additional strategies. What strategy subprofiles may i 's opponents play? First of all, i is sure that the only possible candidates are those in the set

$$\mathbb{B}_i(-i) \triangleq \{\theta_{-i} : \theta_{-i} \in \mathbb{B}_i\}.$$

Second of all, since i believes that his opponents are rational, for each specific valuation subprofile $\theta_{-i} \in \mathbb{B}_i(-i)$, assuming that $\theta_{-i} = \theta_{-i}^*$, i is sure that the strategy subprofiles that $-i$ might consider playing are those in

$$U_{-i}^i(\theta_{-i}) \triangleq \prod_{j \in -i} U_j(\theta_j, S).$$

Third and finally, since he cannot tell which subprofile in $\mathbb{B}_i(j)$ is the true one, i can only be sure that the strategy subprofile that will actually be played by his rational opponents is

$$U_{-i}^i \triangleq \bigcup_{\theta_{-i} \in \mathbb{B}_i(-i)} U_{-i}^i(\theta_{-i}).$$

Accordingly, i will conservatively eliminate a strategy $s_i \in U_i(\theta_i^*, S)$ if there exists an alternative mixed strategy $\sigma_i \in \Delta(U_i(\theta_i^*, S))$, such that, for *all* strategy subprofiles $s_{-i} \in U_{-i}^i$, s_i is strictly dominated by s_i' (relative to θ_i^* and U).

Definition 5. *The set of strictly conservative strategies of player i is*

$$\mathcal{C}_i \triangleq U_i(\theta_i^*, S) \setminus \{s_i : \exists \sigma_i \in \Delta(U_i(\theta_i^*, S)) \text{ such that } s_i <_{U_i \times U_{-i}^i}^{\theta_i^*} \sigma_i\}.$$

The set of all profiles of strictly conservative strategies is $\mathcal{C} = \mathcal{C}_1 \times \dots \times \mathcal{C}_n$.

4 Our mechanism For the unit-demand Case

Let us present our result for the simpler, unit-demand setting.

Definition 6. *A multi-unit auction is unit-demand if, for all players i , $v_i(2) = \dots = v_i(m) = 0$.*

Theorem 1. *There exists a mechanism M_ϵ that (1) is individually rational, (2) never generates negative revenue, and (3) under all profiles of strictly conservative strategies, generates expected revenue*

$$\geq (1 - \epsilon) \left(1 - \frac{\lceil \sqrt{2m} \rceil}{m} \right) CEB .$$

The Interesting Case. The interesting case of Theorem 1 is $m < n$. Indeed, a mechanism much simpler than M_ε virtually achieves revenue at least CEB when $m \geq n$. Accordingly, the reader may assume that $m < n$ in order to justify dealing with some of M_ε 's complications and losing a factor $\frac{\lceil \sqrt{2m} \rceil}{m} CEB$ in revenue.

4.1 The Intuition Behind M_ε

We first describe M_ε assuming that the players consider only strategies that survive two levels of elimination of *weakly* dominated strategies, relative to the players' beliefs. (Recall that, essentially, a strategy s_i is weakly dominated by another—possibly mixed—strategy σ_i of i , if σ_i provides i with a utility that is at least as large as that provided by s_i , no matter what strategies i 's opponents may use, and strictly greater than that provided by s_i for at least some strategies of i 's opponents.)

Mechanism M_ε is of normal form. Initially, each player i , simultaneously with his opponents, reports a profile of non-negative real values, $v_i = (v_i(1), \dots, v_i(n))$, where $v_i(j)$ is i 's *report about j* .

The mechanism wishes to incentivize i so that $v_i(j)$ is greater than or equal to the highest value v for which, according to his beliefs, i is sure that $\theta_j^* \geq v$, that is, $v_i(j) \geq \min_{\theta \in \mathbb{B}_i} \theta_j$. In particular, therefore, M_ε wishes to ensure that $v_i(i) \geq \theta_i^*$.

First, let us describe a simple attempt to achieve the above desiderata that does not work, and then present a modification that does.

The attempt is as follows. First of all, for each value $v_i(j)$ reported by player i about an opponent j , the mechanism rewards i with an amount of money $r(v_i(j))$, where the reward function r is strictly increasing and such that the total reward given by M_ε is upperbounded by εCEB . We stress that M_ε does not offer rewards for i 's “self report”, that is, $v_i(i)$.

Next, for each player j , the mechanism sets $K_{MAX}(j)$ to be the maximum value reported about j by another player. If the players reported truthfully, then $K_{MAX}(j)$ represents the best “guaranteed” *knowledge* that the other players have about j 's valuation. The mechanism also sets $C_{MAX}(j)$ to be the (lexicographically first) player who reports $K_{MAX}(j)$ about j . One can think of $C_{MAX}(j)$ as the opponent of j who is most *confident* about j 's valuation.

Next, the mechanism sorts the players, in decreasing order, according to their K_{MAX} values. That is, it computes a permutation of the players, p_1, \dots, p_n , such that $K_{MAX}(p_i) \geq K_{MAX}(p_j)$ whenever $i < j$. Then, for the first m players, that is, for each of p_1, \dots, p_m , M_ε does the following. It compares $K_{MAX}(p_i)$ with $v_{p_i}(p_i)$. If $K_{MAX}(p_i) \leq v_{p_i}(p_i)$, then p_i gets a copy of the good and pays $K_{MAX}(p_i)$ for it. Otherwise, $C_{MAX}(p_i)$ has to pay $K_{MAX}(p_i)$, and no copy is allocated in this “transaction”.

So far, it should be clear that,

Under-report one's own valuation is weakly dominated.

Indeed, for a player i , the strategy consisting of reporting a profile of values v_i such that $v_i(i) \leq \theta_i^*$ is weakly dominated by reporting the profile v_i^* such that

$$v_i^*(i) = \theta_i^* \quad \text{and} \quad v_i^*(j) = v_i(j) \text{ for all } j \neq i .$$

Note that, due to the rewards offered by the mechanism, it may sometimes be rational for a player i to report a value $v_i(j)$ that he knows to be greater than $\max_{\theta \in \mathbb{B}_i} \theta_j$. However, such “over-reporting” is *not* a problem, because it is easy to see that it can only but *increase* the total revenue generated.

The problem with the present mechanism M_ε is that

Under-reporting one’s beliefs about his opponents’ valuations need not be weakly dominated.

This is a problem, because M_ε generates revenue only based on the values the players report about their opponents. However, it may not be weakly dominated for i to report $v_i(j) = 0$ for each opponent i , despite the fact that i believes that j ’s guaranteed value for a copy of the good is positive, or even very high. Let us now give an example of why this may happen.

Example Let t_i , the “truthful” report of i , be so defined:

$$t_i(i) = \theta_i^* \text{ and } t_i(j) = \min_{\theta \in \mathbb{B}_i} \theta_j \text{ for all } j \neq i.$$

Let us now describe some additional beliefs that may cause reporting $v_i(j) = 0$ for all $j \neq i$ to be not weakly dominated, despite the fact that $t_i(j)$ is positive.

We wish to stress that the beliefs we are about to describe may be “far fetched”, but do *not* “contradict \mathbb{B}_i ”, and thus should *not* interfere with the revenue performance of the mechanism.

Suppose that player i believes that his opponents will report a value subprofile v_{-i} such that, for a specific player $j \neq i$, (a) $v_k(j) = 0$ for all $k \neq i, j$ and (b) under the report (t_i, v_{-i}) :

$$i = p_{m+1}, \quad j = p_\ell \text{ for } \ell \leq m, \quad 0 < K_{MAX}(i) < t_i(j) \ll \theta_i^*.$$

These additional beliefs of i have three consequences. First, $C_{MAX}(j) = i$. In fact, all other players report 0 about j , while i reports truthfully and $t_i(j)$ is positive. Second, i believes that, by bidding truthfully in mechanism M_ε , he will not get any copy, and his only utility will come from the small rewards he receives for the values he reports about his opponents. Third, i believes that if he reports the valuation profile t'_i that coincides with the truthful report t_i except for component j , where $t'_i(j) = 0$, then he will get a copy of the good (at a very good price), and thus, despite that he loses the small reward $r(t_i(j))$, he will receive a much higher utility than that from the rewards alone. Accordingly, i is better off reporting 0 about j instead of his true belief about j . \triangle

To avoid this unfortunate scenario, we modify the mechanism as follows. After receiving the players’ reports, giving them their rewards, and computing the reordering p_1, \dots, p_n , the mechanism tries to “sell” a copy of the good as before (and thus at the same price as before) to each player $p_i \in W$, where $W = \{p_1, \dots, p_{m - \lceil \sqrt{2m} \rceil}\}$.

At this point, M_ε also tries to sell a copy of the good, as before, to each player $p_i \notin W$ such that

- (\star) $i = \lceil \sqrt{2m} \rceil + k$ (where $k \geq 1$) and there exist at least k players p_j , $j < i$, for which $C_{MAX}(p_j) = p_i$.

Let us explain. Conceptually, the mechanism sets aside $\lceil\sqrt{2m}\rceil$ copies of the good. Each of the remaining $m - \lceil\sqrt{2m}\rceil$ copies is “offered” to a player p_j in W at price $K_{MAX}(p_j)$, and the copy is allocated if the offer is “accepted”, that is, if $K_{MAX}(p_j) \leq v_{p_j}(p_j)$ —i.e., if the highest report about p_j of p_j ’s opponents is at most p_j ’s self-report. Thus, while the players in W are guaranteed to be offered a copy, those outside W are not, and thus stand to lose a possibly high utility. It is thus not surprising that every player i would prefer to belong to W . Notice that the cardinality of W is fixed to be $m - \lceil\sqrt{2m}\rceil$, and recall that a player j ends up in W if the highest report about j is sufficiently high. Thus a player i may worry that, reporting a high value for one or more opponents j may result in “kicking himself out of W ”. Accordingly, if the mechanism does not take any counter-measures, it risks that a player i may under-report about his opponents. To prevent this, the mechanism essentially enters the following “contract” with any player:

“Dear player i : if you are outside W , but would have been in W if you reported 0 about all of your opponents, then I will offer you a copy of the good anyway, at price $K_{MAX}(i)$ ”.

Notice that this contract is technically expressed by the above condition (\star). Under such condition, it becomes intuitive that player i would prefer to report about his opponents, and collect the rewards that otherwise would have forfeited. This informal intuition could be turned into a formal argument, ultimately showing that mechanism M_ε guarantees revenue $\geq (1 - \varepsilon)(1 - \frac{\lceil\sqrt{2m}\rceil}{m})CEB$, if the players only consider strategies surviving two levels of elimination of weakly dominated strategies (according to their beliefs).

To ensure that such revenue continue to be guaranteed under our more demanding solution concepts (i.e., in strictly conservative strategies) we need to modify M_ε . In our modification, no cash rewards will be given for any reported value. However, by reporting $v_i(j)$, player i has a probability of getting a copy of the good for free, where this probability, although strongly upperbounded, is also strictly increasing with the reported $v_i(j)$. These “alternative rewards” are actually offered also for self-reports, that is, for the reported values $v_i(i)$.

The details of this modification can be found in our more formal description of M_ε of Section 4.2, and an explanation of why they work can be found in Section 4.3. Here we only wish to point out why it was not a good idea to offer cash rewards for self-reports, but it is OK to provide “alternative rewards” for self-reports.

If the mechanism provides a positive cash reward for each positive reported value, including a positive self-report, then the generated revenue might be negative. For instance, this happens when some self-reported value is positive and all other reported values are 0. (Notice that this might happen when the players have no information about their opponents.) By contrast, with our “alternative rewards”, the mechanism never generates negative revenue. This is because we do not pay the rewards in cash. In particular, as we shall see, in the above scenario, a player i self-reporting a positive $v_i(i)$ expects to receive a good for free with probability $\frac{1}{n^2} \cdot \frac{1}{1+v_i(i)}$, and the revenue generated is 0.

4.2 The Mechanism M_ε

In M_ε , for each player i , P_i is a real-valued variable, initially set to 0, whose final value will be the price of i ; and, for each player i , A_i is a binary variable, initially set to 0, whose final value is 1 if and only if i is assigned a copy of the good. All sentences between quotation marks are comments.

We find it simple to break down the mechanism into 5 conceptual steps, marked a to e .

MECHANISM M_ϵ

Each player i , simultaneously with his opponents, reports a profile v_i of non-negative reals.

“Allegedly $v_i(j) = \min_{\theta \in \mathbb{B}_i} \theta_j$ for all opponents j of i .”

a: With probability ϵ , choose uniformly at random a pair of players $(i, j) \in \{1, \dots, n\} \times \{1, \dots, n\}$; set $A_i = 1$ with probability $1 - \frac{1}{1+v_i(j)}$; and HALT.

“ $A_j = 0$ for $j \neq i$ and $P_j = 0$ for all all j .”

With complementary probability, execute the following 4 Steps.

b: For each player j , set $K_{MAX}(j) = \max_{i \neq j} v_i(j)$ and $C_{MAX}(j) = \operatorname{argmax}_{i \neq j} v_i(j)$, with ties broken lexicographically.

c: Compute the permutation P of the players, p_1, \dots, p_n , such that $K_{MAX}(p_1) \geq K_{MAX}(p_2) \geq \dots$, where ties are broken lexicographically.

d: Let $W = \{p_1, \dots, p_{m - \lceil \sqrt{2m} \rceil}\}$.

For each $p_i \in W$, “ p_i gets an offer for a copy at price $K_{MAX}(p_i)$ ”:

If $K_{MAX}(p_i) \leq v_{p_i}(p_i)$ then $A_{p_i} := 1$ and $P_{p_i} := P_{p_i} + K_{MAX}(p_i)$. “ p_i gets, and pays for, a copy”

Else, $P_{C_{MAX}(p_i)} := P_{C_{MAX}(p_i)} + K_{MAX}(p_i)$. “ $C_{MAX}(i)$ is penalized for his bad report.”

e: For each $p_i \notin W$, let C_{p_i} be the set of all players p_j , with $j < i$, such that $C_{MAX}(p_j) = p_i$. For each $p_i \notin W$, such that $i = \lceil \sqrt{2m} \rceil + k$ (for $k \geq 1$) and $|C_{p_i}| \geq k$, p_i gets an offer for a copy at price $K_{MAX}(p_i)$ as follows:

If $K_{MAX}(p_i) \leq v_{p_i}(p_i)$ then set $A_i := 1$ and $P_{p_i} := P_{p_i} + K_{MAX}(p_i)$.

Else set $P_{C_{MAX}(p_i)} := P_{C_{MAX}(p_i)} + K_{MAX}(p_i)$.

4.3 Proof of Theorem 1

Lemma 0 M_ϵ allocates at most $m - \lceil \sqrt{2m} \rceil$ copies in Step d and at most $\lceil \sqrt{2m} \rceil$ copies in Step e.

Proof. It is clear that M_ϵ allocates, in step d, at most $m - \lceil \sqrt{2m} \rceil$ copies to the players in W . It is thus enough to prove that M_ϵ will allocate, in step e, at most $\lceil \sqrt{2m} \rceil$ copies to the players outside W . We will now prove this.

Let the players that receive a copy in step e be p_{x_1}, \dots, p_{x_y} , where $m - \lceil \sqrt{2m} \rceil < x_1 < \dots < x_y$. Then, for each such player p_{x_i} , we have $|C_{p_{x_i}}| \geq x_i - (m - \lceil \sqrt{2m} \rceil)$. It is clear that for $i \neq j$, $C_{p_i} \cap C_{p_j} = \emptyset$. Moreover, for all $i \leq y$, $C_{p_{x_i}} \subset \{p_1, p_2, \dots, p_{x_y}\}$. Thus, $\bigcup_{1 \leq i \leq y} C_{p_{x_i}} \subset \{p_1, \dots, p_{x_y}\}$. This implies that

$$\sum_{i=1}^y |C_{p_{x_i}}| = \left| \bigcup_{1 \leq i \leq y} C_{p_{x_i}} \right| \leq |\{p_1, \dots, p_{x_y}\}| = x_y.$$

We also have

$$\sum_{i=1}^y |C_{p_{x_i}}| = \sum_{i=1}^y (x_i - (m - \lceil \sqrt{2m} \rceil)) = \sum_{i=1}^y x_i - \sum_{i=1}^y (m - \lceil \sqrt{2m} \rceil) \geq$$

$$x_y + \sum_{j=1}^{y-1} (m - \lceil \sqrt{2m} \rceil + j) - \sum_{j=1}^y (m - \lceil \sqrt{2m} \rceil) = x_y - (m - \lceil \sqrt{2m} \rceil) + \sum_{j=1}^{y-1} j > x_y - m + \frac{y(y-1)}{2}.$$

In sum, $\sum_{i=1}^y |C_{p_{x_i}}| > x_y - m + \frac{y(y-1)}{2}$. Accordingly, since $\sum_{i=1}^y |C_{p_{x_i}}| \leq x_y$, we have

$$x_y - m + \frac{y(y-1)}{2} < x_y \Rightarrow y(y-1) < 2m \Rightarrow y < \sqrt{2m} + 1 \Rightarrow y < \lceil \sqrt{2m} \rceil. \quad \blacksquare$$

Lemma 1 M_ε is individually rational.

Proof. It is easy to see that, if a player i reports $v_i(j) = 0$ for all opponents j , then his utility is at least 0 no matter what his opponents might report. \blacksquare

Recall that, in M_ε , for each player i , $S_i = (\mathbb{R}_{\geq 0})^n$, and that $S = S_1 \times \dots \times S_n$.

Lemma 2 $\forall i \ U_i(\theta_i^*, S) \subset \{v_i \in (R_{\geq 0})^n : v_i(i) \geq \theta_i^*\}$.⁴

Proof. Let $v_i(i) < \theta_i^*$. Define the strategy v_i^* as follows.

$$v_i^*(j) = \begin{cases} v_i(j) & \text{if } j \neq i \\ \theta_i^* & \text{otherwise.} \end{cases}$$

We now prove that v_i is strictly dominated by v_i^* : more precisely, $v_i <_S^{\theta_i^*} v_i^*$. That is, we prove:

$$u_i(\theta_i^*, (v_i, s_{-i})) < u_i(\theta_i^*, (v_i^*, s_{-i})) \text{ for all } s_{-i} \in S_{-i}.$$

Let s_{-i} be an arbitrary strategy subprofile in S_{-i} and let us compare the two utilities of i when his opponents report s_{-i} and he reports either v_i or v_i^* : that is, $u_i(\theta_i^*, (v_i, s_{-i}))$ and $u_i(\theta_i^*, (v_i^*, s_{-i}))$.

Both utilities have three components: a “reward component”, a “penalty component”, and a “copy component”.

Player i 's reward component is the *expected* utility that he receives in step a . Thus, no matter what the strategy subprofiles of i 's opponents are, i 's reward components for his strategies v_i and v_i^* respectively are

$$r(v_i) = \sum_j \varepsilon \cdot \frac{1}{n^2} \cdot \left(1 - \frac{1}{1 + v_i(j)}\right) \cdot \theta_i^* \quad \text{and} \quad r(v_i^*) = \sum_j \varepsilon \cdot \frac{1}{n^2} \cdot \left(1 - \frac{1}{1 + v_i^*(j)}\right) \cdot \theta_i^*.$$

In fact, in step a , the mechanism chooses the pair (i, j) with probability ε/n^2 , and then assigns a free copy of the good to i with probability $\varepsilon \frac{1}{1+v_i(j)}$, if i reports v_i , and with probability $\varepsilon \frac{1}{1+v_i^*(j)}$, if i reports v_i^* . Moreover, i 's utility when he receives a copy of the good for free always is θ_i^* , no matter what he reports.

⁴Actually, $U_i(\theta_i^*, S_i) = \overline{\{v_i \in (R_{\geq 0})^n : v_i(i) \geq \theta_i^*\}}$. Proving equality of the two sets is not hard, but is not ultimately required to prove Theorem 1, which is our main interest.

Since $v_i(j) = v_i^*(j)$ when $j \neq i$, and also $0 \leq v_i(i) < \theta_i^*$, we have that:

$$r(v_i) - r(v_i^*) = \frac{\varepsilon}{n^2} \cdot \left(1 - \frac{1}{(1 + v_i(i))}\right) - \frac{\varepsilon}{n^2} \cdot \left(1 - \frac{1}{(1 + \theta_i^*)}\right) = \frac{\varepsilon}{n^2} \cdot \left(\frac{v_i(i) - \theta_i^*}{(1 + v_i(i))(1 + \theta_i^*)}\right) < 0,$$

That is, $r(v_i) < r(v_i^*)$. In sum, v_i^* strictly dominates v_i with respect to i 's reward component.

Let us now consider i 's penalty component, assuming that M_ε does not halt in Step a , because otherwise the penalty component is 0. This component solely depends on (a) i 's reports about his opponents, and (b) i 's opponents self-reports. (In fact, these reports fully determine the values $K_{MAX}(j)$ for each player j , and thus also the player permutation p_1, \dots, p_n , and finally also the players $C_{MAX}(p_j)$ for all j . Thus, i pays a penalty $K_{MAX}(p_j)$, in Step d , for some player p_j if and only if (a) $C_{MAX}(p_j) = i$ and (b) what p_j reports about himself, $s_j(j)$, is less than what i reports about j , $K_{MAX}(p_j)$.) Since these reports are identical under the strategy profiles (v_i, s_{-i}) and (v_i^*, s_{-i}) , i 's penalty component must also be identical under the two profiles.

Finally, let us consider the copy component. This component consists of the utility i gets when he is allocated a copy of the good in Step d or e .

Again assuming that M_ε does not halt in Step a , let us compare the copy component of i 's utility under the strategy profiles (v_i, s_{-i}) and (v_i^*, s_{-i}) . There are three cases to consider: namely,

- (1) i is not allocated any copy under (v_i, s_{-i}) and (v_i^*, s_{-i}) ;
- (2) i is allocated a copy both under (v_i, s_{-i}) and (v_i^*, s_{-i}) ; and
- (3) i is allocated a copy under (v_i^*, s_{-i}) , but not under (v_i, s_{-i}) .⁵

In case 1, i 's copy utility is the same under (v_i, s_{-i}) and (v_i^*, s_{-i}) , because he always pays 0 for any unallocated copy anyway. In case 2, i 's copy utility is also the same under (v_i, s_{-i}) and (v_i^*, s_{-i}) , because his valuation for the copy received always is θ_i^* (independently of what he reports about himself) and the price he pays is $K_{MAX}(i)$, which as argued above solely depends on i 's opponents' reports. In case 3, i 's copy utility is 0 under profile (v_i, s_{-i}) , and consists of the non-negative number $\theta_i^* - K_{MAX}(i)$ under (v_i^*, s_{-i}) .

In sum, the first component of i 's utility is strictly less under the strategy profile (v_i, s_{-i}) than under the strategy profile (v_i^*, s_{-i}) , while all the other components of i 's utility are no greater under (v_i, s_{-i}) than under (v_i^*, s_{-i}) . Thus v_i is strictly dominated by v_i^* . ■

Lemma 3 $\mathcal{C}_i \subset \{v_i \in U_i(\theta_i^*, S) : v_i(j) \geq \min_{\theta \in \mathbb{B}_i} \theta_j \text{ for all } j \neq i\}$

Proof. Fix $v_i \in \mathcal{C}_i$ and $j \neq i$. We will argue by contradiction $v_i(j) \geq \min_{\theta \in \mathbb{B}_i} \theta_j$. Define

$$m_i(j) \triangleq \min_{\theta \in \mathbb{B}_i} \theta_j.$$

Assume that $v_i(j) < m_i(j)$ and define the strategy v_i^* as follows.

$$v_i^*(j) = \begin{cases} v_i(k) & \text{if } k \neq j \\ m_i(j) & \text{if } k = j. \end{cases}$$

⁵The case where i is allocated a copy under (v_i, s_{-i}) , but not under (v_i^*, s_{-i}) cannot occur, because $v_i(i) < v_i^*(i)$.

We reach a contradiction by proving that that

$$v_i <_{U_i \times U_{-i}}^{\theta_i^*} v_i^* .$$

Recall that $U_i \triangleq U_i(\theta^*, S)$, that $U_{-i}^i(\theta_{-i}) \triangleq \prod_{j \in -i} U_j(\theta_j, S)$, and that $U_{-i}^i \triangleq \bigcup_{\theta_{-i} \in \mathbb{B}_i(-i)} U_{-i}^i(\theta_{-i})$.

Thus we reach a contradiction by proving that

$$u_i(\theta_i^*, (v_i, s_{-i})) < u_i(\theta_i^*, (v_i^*, s_{-i})) \quad \text{for all } s_{-i} \in \overline{U}_i .$$

Let s_{-i} be an arbitrary strategy subprofile in \overline{U}_i , and let us compare the two utilities of i under the strategy profiles $(\theta_i^*, (v_i, s_{-i}))$ and $(\theta_i^*, (v_i^*, s_{-i}))$. Again, we break our analysis under the three components of i 's utility: reward, penalty, and copy. The analysis of the reward component is essentially the same as in the proof of Lemma 2, but that of the penalty and the copy component is different.

ANALYSIS OF THE REWARD COMPONENT. Player i 's reward components (again, i 's *expected* utility in step a) under the above two strategy profiles respectively are

$$r(v_i) = \sum_k \varepsilon \cdot \frac{1}{n^2} \cdot \left(1 - \frac{1}{1 + v_i(k)}\right) \cdot \theta_i^* \quad \text{and} \quad r(v_i^*) = \sum_k \varepsilon \cdot \frac{1}{n^2} \cdot \left(1 - \frac{1}{1 + v_i^*(k)}\right) \cdot \theta_i^* .$$

Since $v_i(k) = v_i^*(k)$ when $k \neq j$, and also $0 \leq v_i(j) < m_i(j)$, we have that:

$$r(v_i) - r(v_i^*) = \frac{\varepsilon}{n^2} \cdot \left(1 - \frac{1}{(1 + v_i(j))}\right) - \frac{\varepsilon}{n^2} \cdot \left(1 - \frac{1}{(1 + m_i(j))}\right) = \frac{\varepsilon}{n^2} \cdot \left(\frac{v_i(j) - m_i(j)}{(1 + v_i(i))(1 + \theta_i^*)}\right) < 0,$$

That is, $r(v_i) < r(v_i^*)$. In sum, v_i^* strictly dominates v_i with respect to i 's reward component.

ANALYSIS OF THE PENALTY COMPONENT. Assume that M_ε does not halt in Step a , because otherwise the penalty component is 0. Again, i 's penalty component solely depends on (a) i 's reports about his opponents, and (b) i 's opponents self-reports.

Notice that it suffices to prove that

$$\begin{aligned} & \text{if } i \text{ gets a penalty } v_i(k) \text{ under profile } (\theta_i^*, (v_i^*, s_{-i})), \\ & \text{then he gets the same penalty under } (\theta_i^*, (v_i, s_{-i})). \end{aligned}$$

First of all, notice that, under $(\theta_i^*, (v_i^*, s_{-i}))$, i believes that he will not get penalty $m_i(j) = \min_{\theta \in \mathbb{B}_i} \theta_j$ for reporting $v_i^*(j) = m_i(j)$. In fact, i believes that j is rational, and thus that he will not play a strictly dominated strategy. Thus, by Lemma 2, $s_j(j) \geq m_i(j)$.

Assume now that i gets a penalty $v_i^*(k)$ for some $k \neq j$. Recall that, for i to get such a penalty, it is necessary that, in Step d or e , player k gets an offer to buy a copy of the good at price $v_i(k)$, and that $i = C_{MAX}(k)$. Also recall that $W = \{p_1, \dots, p_{m - \lfloor \sqrt{2m} \rfloor}\}$. We now distinguish 2 cases.

1. $k \in W$ under profile $(\theta_i^*, (v_i^*, s_{-i}))$.

Let W_k^* be the set of players that, under profile $(\theta_i^*, (v_i^*, s_{-i}))$, precede k in permutation P^* , computed in Step 3 of the mechanism. That is, W_k^* consists

of all players ℓ such that $K_{MAX}(\ell) > K_{MAX}(k)$ and all players m that are lexicographically smaller than k and such that $K_{MAX}(m) = K_{MAX}(k)$.

Similarly define W_k for the profile $(\theta_i^*, (v_i, s_{-i}))$. Under this profile, for every $\ell \neq j$, $K_{MAX}(\ell)$ has the same value as under profile $(\theta_i^*, (v_i^*, s_{-i}))$. Moreover, the value of $K_{MAX}(j)$ under profile $(\theta_i, (v_i, s_{-i}))$ is at most the value of $K_{MAX}(j)$ under profile $(\theta_i^*, (v_i, s_{-i}))$. So, $W^k \subset W_k^*$.

In sum, k gets the same offer under both profiles, and i pays the same penalty $v_i(k)$.

2. $k \notin W$ under profile $(\theta_i^*, (v_i^*, s_{-i}))$.

Let k 's position in permutation P be n , or equivalently $p_n = k$. Define W_k and W_k^* as in case 1. Since i gets a penalty $v_i(k)$ under $(\theta_i^*, (v_i^*, s_{-i}))$, we have that W_k^* contains at least $n - (m - \lceil \sqrt{2m} \rceil)$ players whose C_{MAX} is k . Clearly, these $n - (m - \lceil \sqrt{2m} \rceil)$ players are also in W_k , because only the value $K_{MAX}(k)$ is smaller under $(\theta_i, (v_i, s_{-i}))$, while the remaining K_{MAX} values are the same as under $(\theta_i^*, (v_i, s_{-i}))$. Moreover, by exactly the same argument as in case 1, we have that $W_k \subset W_k^*$, which implies $|W_k| \leq |W_k^*|$.

So, i also gets the penalty $v_i(k)$ under profile $(\theta_i^*, (v_i^*, s_{-i}))$.

ANALYSIS OF THE COPY COMPONENT. Assume again M_ϵ does not halt in Step *a*. The copy component consists of the utility i gets when he is allocated a copy of the good in Step *d* or *e*. Whether player i gets a copy and at what price depends only on i 's self report and on the reports of i 's opponents. Since these are the same under both profiles $(\theta_i, (v_i, s_{-i}))$ and $(\theta_i^*, (v_i^*, s_{-i}))$, the copy component is the same under both profiles.

In sum, the first component of i 's utility is strictly less under the strategy profile (v_i, s_{-i}) than under the strategy profile (v_i^*, s_{-i}) , while all the other components of i 's utility are no greater under (v_i, s_{-i}) than under (v_i^*, s_{-i}) . Thus v_i is strictly dominated by v_i^* , more precisely $v_i < \frac{\theta_i^*}{U_i} v_i^*$. ■

Finishing the Proof of Theorem 1

Mechanism M_ϵ is individually rational by Lemma 1 and never generates negative revenue, because it never transfers any money to a player. Finally, let us analyze its expected revenue.

Recall that, for each player j , we have $K_{MAX}(j) = \max_{i \neq j} s_i(j)$. Similarly, for each player j , define

$$M_{MAX}(j) = \max_{i \neq j} m_i(j).$$

By Lemmas 2 and 3, for all profiles s of strictly conservative strategies, $s_i(i) \geq \theta_i^*$, and $s_i(j) \geq m_i(j)$ for all $j \neq i$. Thus, $K_{MAX}(j) \geq M_{MAX}(j)$ for all players j .

Now recall that, if Q is a permutation (q_1, \dots, q_n) of the n players such that $M_{MAX}(q_1) \geq \dots \geq M_{MAX}(q_n)$, then $CEB = \sum_{i=1}^m M_{MAX}(q_i)$. Thus, $\sum_{i=1}^m K_{MAX}(q_i) \geq CEB$.

Now notice that M_ε actually computes a permutation of the players $P = (p_1, p_2, \dots, p_n)$ such that $K_{MAX}(p_1) \geq \dots \geq K_{MAX}(p_n)$. Thus

$$\sum_{i=1}^m K_{MAX}(p_i) \geq \sum_{i=1}^m K_{MAX}(q_i) \geq CEB$$

At the same time, for each $i \in \{1, \dots, m - \lceil \sqrt{2m} \rceil\}$, our mechanism offers to sell a copy to player p_i at price $K_{MAX}(p_i)$. If $v_{p_i}(p_i) \geq K_{MAX}(p_i)$, then p_i pays $K_{MAX}(p_i)$ to M_ε . Otherwise, $C_{MAX}(p_i)$ pays $K_{MAX}(p_i)$ to M_ε as a penalty. So M_ε guarantees revenue at least $\sum_{i=1}^{m - \lceil \sqrt{2m} \rceil} K_{MAX}(p_i)$.

Let us now show the following:

Claim 1. For all decreasing sequences $v_1 \geq v_2 \geq \dots$ and positive integers x and k , we have

$$(x+k) \sum_{i=1}^x v_i \geq x \sum_{i=1}^{x+k} v_i .$$

Proof of Claim 1.

$$\begin{aligned} (x+k) \sum_{i=1}^x v_i &= x \sum_{i=1}^x v_i + k \sum_{i=1}^x v_i \geq x \sum_{i=1}^x v_i + k \sum_{i=1}^x v_x = x \sum_{i=1}^x v_i + kxv_x \geq \\ &x \sum_{i=1}^x v_i + xkv_{x+1} = x \sum_{i=1}^x v_i + x \sum_{i=x+1}^{x+k} v_{x+1} \geq x \sum_{i=1}^x v_i + x \sum_{i=x+1}^{x+k} v_i = x \sum_{i=1}^{x+k} v_i . \end{aligned}$$

□

Claim 1 can be restated as

$$\sum_{i=1}^x v_i \geq \frac{x}{x+k} \sum_{i=1}^{x+k} v_i .$$

Thus letting $v_i = K_{MAX}(i)$, $x = m - \lceil \sqrt{2m} \rceil$, and $k = \lceil \sqrt{2m} \rceil$, we have

$$\sum_{i=1}^{m - \lceil \sqrt{2m} \rceil} K_{MAX}(p_i) \geq \frac{(m - \lceil \sqrt{2m} \rceil)}{m} \sum_{i=1}^m K_{MAX}(p_i) \geq \left(1 - \frac{\lceil \sqrt{2m} \rceil}{m}\right) CEB .$$

This ends the proof of Theorem 1. ■

5 Our mechanism for the q -unit demand case

We now present our result for the q -unit demand case.

Definition 7. A multi-unit auction is q -unit demand if, for all players i , $v_i(q+1) = \dots = v_i(m) = 0$.

Theorem 2. M_ε^q is individually rational, it never generates negative revenue, and its expected revenue is at least $(1 - \varepsilon) \left(1 - \frac{q(\lceil \sqrt{\frac{2m}{q}} \rceil + 1)}{m}\right) CEB$.

The Interesting Case: $m < qn$. A much simpler version of M_ε^q virtually achieves revenue at least CEB when $m \geq qn$. Accordingly, one may assume that $m < qn$ when analyzing M_ε^q .

Quick summary of M_ε^q

In addition to the previously used allocation variables and price variables A_i and P_i , M_ε^q also employs, for each player i , a binary variable B_i , that, before the mechanism starts, is set to 0.

The mechanism starts with all players allegedly revealing their beliefs. That is, each player i reports, for each player j and copy c , a non-negative value, $v_i(j)(c)$, which is supposedly the highest value v such that i is sure that $\theta_j^*(c) \geq v$.

The mechanism orders these reported n^2q values in decreasing order, breaking ties appropriately. (For simplicity, assume that there are no ties.)

Similarly to mechanism M_ε , M_ε^q conceptually sets apart $q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right)$ copies of the good. Then, in order, it processes the first $m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right)$ reported values as follows.

When it is the turn of value $v_i(j)(c)$, the mechanism simulates player i offering a copy of the good to player j for price $v_i(j)(c)$. To evaluate whether j accepts or rejects the offer, the mechanism uses the corresponding self-report of j , that is, $v_j(j)(c)$. If j accepts i 's offer, then the mechanism allocates a copy of the good to j and collects from him a payment of $v_i(j)(c)$. If j rejects i 's offer, then the mechanism imposes a penalty of $v_j(j)(c)$ to player i , for his bad report, and irrevocably sets the variable B_j to 1. By doing so, the mechanism remembers never to allocate another copy of the good to j , whether or not j will receive another offer. Furthermore, if j indeed receives such an offer from some player i' , not only will the offer be rejected automatically, but i (rather than i') will be imposed the price offered by i' as an additional penalty. (This behavior of M_ε^q is made necessary by a circumstance that does not arise in the unit-demand case. Assume that player j values a first copy of the good at 100 and a second one at 10, and that he believes that he will be offered a first copy for 90 and a second copy for 9. Then, if he truthfully reports his own valuation, his total utility would be $110 - 99 = 11$. However, if he could reject the first offer and accept only the second one, then he would receive a single copy of the good and his utility would be $100 - 9 = 91$. It is to avoid under-reporting one's true valuation that M_ε^q , once a player rejects an offer, forces him to reject all future ones. At the same time, in order to achieve its revenue target, M_ε^q must be compensated for all rejected offers. But for automatically rejected offers, it cannot penalize the players making them, because their offers might have been conservatively chosen based on their beliefs. It thus penalize the player making the first rejected offer to a player j for all the offers that j automatically rejects.)

After all first $m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right)$ values have been processed, the mechanism M_ε^q , similarly to M_ε , tries to allocate the $q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right)$ copies of the good that have been set aside, by simulating offers using, in order, to the yet unprocessed $n^2q - \left(m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right) \right)$ reported values.

When it is the turn of value $v_i(j)(c)$, an offer to buy a copy of the good for price $v_i(j)(c)$ is made to player j if and only if he was responsible for making at least k of the offers considered so far, where k is the position in which value $v_i(j)(c)$ appears in the ordering.

Again, players who “over-report” their beliefs may be in trouble in M_ε^q (and in fact they might be penalized more than in M_ε), but every one is quite safe in reporting, for every opponent j and copy c , $\min_{\theta \in \mathbb{B}_i} \theta_j^*(c)$. The analysis is essentially the same as before. In particular, it shows that under-reporting one’s true value for each copy of the good is strictly dominated by reporting it truthfully. This is so because, similarly to mechanism M_ε , M_ε^q with probability ε chooses at a random a player i , a player j and a copy c , and gives for free a copy of the good to player i with probability $1 - \frac{1}{1+v_i(j)(c)}$, without allocating any additional copies and without any payments.

In mechanism M_ε^q , at most m offers are ever considered. (This implies that at most m copies of the good need to be allocated, and thus that there are sufficiently many copies set aside for secondary allocations.) The same property held for mechanism M_ε , and was in fact proved in Lemma 0. For mechanism M_ε^q , it is proved in (a corresponding) Lemma 0’.

While the rest of the analysis of M_ε^q is a straightforward adaptation of that of M_ε , the proof of Lemma 0’ requires additional work. Accordingly, after presenting the new mechanism, we shall only prove Lemma 0’ and omit the rest of the analysis M_ε^q .

5.1 The Mechanism M_ε^q

As in mechanism M_ε , for each player i , P_i is a real-valued variable, initially set to 0, whose final value will be the price of i , and A_i is an integer valued variable, initially set to 0, whose final value is k if and only if i is assigned k copies of the good.

However, mechanism M_ε^q makes use, for each player i , of an additional binary variable B_i , whose function is described in the following.

We find it simple to break down the mechanism into 5 conceptual steps, marked a to e .

MECHANISM M_ε^q

Each player i , simultaneously with his opponents, reports, for each player j , a vector $v_i(j) \in \mathbb{R}_{\geq 0}^q$ such that $v_i(j)(1) \geq \dots \geq v_i(j)(q) \geq 0$.

“For each j and copy c , $v_i(j)(c)$ allegedly is the highest value v such that i is sure that $\theta_j^*(c) \geq v$.”

a: *With probability ε , choose at random a triple $(i, j, c) \in \{1, \dots, n\} \times \{1, \dots, n\} \times \{1, \dots, q\}$; set $A_i = 1$ with probability $1 - \frac{1}{1+v_i(j)(c)}$; and HALT.*

With complementary probability, execute the following 4 Steps.

b: *For each player j and copy c , set $K_{MAX}(j)(c) = \max_{i \neq j} v_i(j)(c)$ and $C_{MAX}(j)(c) = \operatorname{argmax}_{i \neq j} v_i(j)(c)$, with ties broken lexicographically.*

c: *Compute the permutation P of the player-copy pairs $(i, c) \in \{1, \dots, n\} \times \{1, \dots, q\}$, $P = ((p_1, c_1), \dots, (p_{qn}, c_{qn}))$, such that $K_{MAX}(p_1)(c_1) \geq K_{MAX}(p_2)(c_2) \geq \dots$ and,*

if $K_{MAX}(p_i)(c_i) = K_{MAX}(p_{i+1})(c_{i+1})$, then $c_i \leq c_{i+1}$.

“Ties over the players are broken lexicographically.”

d: For $i = 1, \dots, m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right)$, if $B_{p_i} = 0$, do:

If $K_{MAX}(p_i)(c_i) \leq v_{p_i}(p_i)(c_i)$ then $A_{p_i} := 1$ and $P_{p_i} := P_{p_i} + K_{MAX}(p_i)(c_i)$.

Else, $B_{p_i} = 1$ and $P_{C_{MAX}(p_i)(c_i)} := P_{C_{MAX}(p_i)} + \sum_{c=c_i}^q K_{MAX}(p_i)(c)$

e: Let C_i be the set of all player-copy pairs (p_j, c_j) , with $j < i$, such that $C_{MAX}(p_j)(c_j) = p_i$.

For $i = m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right) + 1, \dots, m$, if $|C_i| \geq i - \left(m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right) \right)$ and $B_{p_i} = 0$, do:

If $K_{MAX}(p_i)(c_i) \leq v_{p_i}(p_i)(c_i)$, then $A_{p_i} := 1$ and $P_{p_i} := P_{p_i} + K_{MAX}(p_i)(c_i)$.

Else, $B_{p_i} = 1$ and $P_{C_{MAX}(p_i)(c_i)} := P_{C_{MAX}(p_i)} + \sum_{c=c_i}^q K_{MAX}(p_i)(c)$.

5.2 Analysis of M_ε^q

Lemma 0' M_ε^q allocates at most $m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right)$ copies in Step d and at most $q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right)$ copies in Step e.

Proof. It is clear that M_ε^q allocates, in step d, at most $m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right)$ copies to the players. Let us thus only prove that M_ε^q will allocate, in step e, at most $q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right)$ copies to the players.

Let the players p_i who receive a c_i -th copy of the good in step e be p_{x_1}, \dots, p_{x_y} , where $m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right) < x_1 < \dots < x_y$. Consider now the vector $(p_{x_1}, \dots, p_{x_y})$ and let z be the number of distinct players that appear in the vector. For each player j among these z players, let $W(j)$ be the greatest integer, $W(j) \leq y$, such that $p_{x_{W(j)}} = j$. Sort the z players in increasing order of W and let the sorted list be j_1, \dots, j_z , with $W(j_i) < W(j_{i+1})$.

For each such player j_i , define q_i to be the number of copies that j_i receives in step e. Given these definitions, we have $|C_{x_{W(j_i)}}| \geq \sum_{k=1}^i q_k$ for $i \in \{1, \dots, z-1\}$ and $|C_{x_{W(j_z)}}| \geq x_{W(j_z)} - \left(m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right) \right)$. It is also clear that for $i \neq \ell$, $C_{x_{W(j_i)}} \cap C_{x_{W(j_\ell)}} = \emptyset$. Moreover, for all $i \leq z$, $C_{x_{W(j_i)}} \subset \{p_1, p_2, \dots, p_{x_{W(j_z)}}\}$. Thus, $\bigcup_{1 \leq i \leq z} C_{x_{j_i}} \subset \{p_1, \dots, p_{x_{W(j_z)}}\}$. This implies that

$$\sum_{i=1}^z |C_{x_{W(j_i)}}| = \left| \bigcup_{1 \leq i \leq z} C_{x_{W(j_i)}} \right| \leq |\{p_1, \dots, p_{x_{W(j_z)}}\}| = x_{W(j_z)}.$$

We also have

$$\sum_{i=1}^z |C_{x_{W(j_i)}}| \geq x_{W(j_z)} - \left(m - q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right) \right) + \sum_{i=1}^{z-1} \binom{i}{k=1} q_i > x_{W(j_z)} - m + \sum_{i=1}^{z-1} ((z-i)q_i)$$

In sum, $\sum_{i=1}^z |C_{p_{x_{W(j_i)}}}| > x_{W(j_z)} - m + \sum_{i=1}^{z-1} ((z-i)q_i)$. Accordingly, since $\sum_{i=1}^z |C_{p_{x_{W(j_i)}}}| \leq x_{W(j_z)}$, we have

$$x_{W(j_z)} - m + \sum_{i=1}^{z-1} ((z-i)q_i) < x_{W(j_z)} \Rightarrow \sum_{i=1}^{z-1} ((z-i)q_i) < m.$$

The number of copies allocated in step e is $\sum_{i=1}^z q_i$. Thus, let us now

$$\text{maximize } \sum_{i=1}^z q_i \text{ subject to}$$

$$a) \sum_{i=1}^{z-1} ((z-i)q_i) < m$$

and $b) v = (q_1, \dots, q_z)$ being a vector of nonnegative integers such that $q_i \leq q$, for all i .

Let $e_i \in \mathbb{R}^n$ be the unit vector whose i_{th} coordinate is 1. Suppose that $v = (q_1, \dots, q_n)$ satisfies the above constraints and maximizes $\sum_{i=1}^z q_i$. For such a vector v , define an update rule to be a procedure that produces a new vector v' as follows:

Choose one pair (k, ℓ) , such that $k < \ell$ and the inequalities $q_k > 0$ and $q_\ell < q$ hold. If such a pair (k, ℓ) exists, set $v' = v - e_k + e_\ell$, otherwise the update is terminated and produces $v' = v$. The vector v' clearly satisfies the constraints and gives $\sum_{i=1}^z q'_i = \sum_{i=1}^z q_i$, so v' also maximizes the desired sum.

Now start with a vector v^0 that maximizes $\sum_{i=1}^z q_i$. At each step, we update vector v^i , getting some vector v^{i+1} , until we reach some v^k such that $v^k = v^{k-1}$ (such a k has to exist, since otherwise each update operation increases by at least 1 the sum in constraint (a)). Set $v^* = v^k$ and recall that v^* also satisfies the constraints and maximizes $\sum_{i=1}^z q_i$. Because v^* cannot be updated, the following must hold: $v_z^* = v_{z-1}^* = \dots = v_{z-k}^* = q$, $0 \leq v_{z-k-1}^* \leq q$, and $v_{z-k-2}^* = \dots = v_1^* = 0$.

From the constraint $\sum_{i=1}^{z-1} ((z-i)v_i^*) < m$, we get $q \frac{k(k-1)}{2} < m \Rightarrow k < \sqrt{\frac{2m}{q}} + 1 \Rightarrow k \leq \left\lceil \sqrt{\frac{2m}{q}} \right\rceil$. So $\sum_{i=1}^z v_i^* = v_{z-k-1}^* + \sum_{i=z-k}^z v_i^* \leq q + qk \leq q \left(\left\lceil \sqrt{\frac{2m}{q}} \right\rceil + 1 \right)$. ■

6 Final Remarks

The Asymptotic Nature of Our Analysis Let us emphasize that the analysis of the revenue performance of our mechanisms M_ε and M_ε^q are asymptotically optimal (as m grows), but are not optimal when m is small. In particular, when $m = 2$, our analysis of M_ε only guarantees revenue 0. However, a more careful analysis for that special case guarantees revenue at least $CEB/2$.⁶

Benchmarking on the Players' Beliefs About Their Opponents We believe that including the players' beliefs will enable new and attractive benchmarks to be defined. To achieve such benchmarks will require new mechanisms and solution concepts. We personally believe that adoption of new solution concept should be welcome, as rationality has many nuances that, though exploitable in many applications of interest, are not captured by traditional solution concepts, such as dominant strategies.

⁶Our general analysis has the mechanism put aside 2 copies when $m = 2$, while one copy is enough. (In particular, in the proof of Lemma 0 we throw away $\left\lceil \sqrt{2m} \right\rceil$ copies, to simplify our analysis.)

Bringing Epistemic Game Theory to Bear Possibilistic beliefs are certainly of theoretical interest, but also of practical interest when one can lowerbound the knowledge the players have about each other. In addition, they usher in the sophisticated machinery of epistemic game theory in mechanism design. We believe and hope that epistemic game theory will prove useful in more (and more complex) applications.

Acknowledgements

We thank Jing Chen for her many insightful comments and her careful reading of this paper.

References

- [1] R. Aumann. Agreeing to Disagree. *Annals of Statistics*, 4: 1236-1239, 1976.
- [2] P. Battigalli and M. Siniscalchi. Rationalization and incomplete information. *Advances of Theoretical Economics*, Volume 3, Issue 1, Article 3, 46 pp., 2003.
- [3] D. Bergemann and S. Morris. Robust mechanism design. *Econometrica*, Vol. 73, No. 6, pp. 1771-1813, 2005.
- [4] D. Bergemann and S. Morris. Robust Mechanism Design: An Introduction. In D. Bergemann and S. Morris, *Robust Mechanism Design*, World Scientific Press, Hackensack, NJ, pp. 1-48, 2012.
- [5] D. Bergemann and S. Morris. Informational Robustness and Solution Concepts. Working paper, 2014.
- [6] J. Chen and S. Micali. Mechanism Design with Set-Theoretic Beliefs. *Journal of Economic Theory*, Vol. 156, pp. 77-102, 2015.
- [7] J. Chen, S. Micali, and R. Pass. Tight Revenue Bounds with Possibilistic Beliefs and Level- k Rationality. *Econometrica*, Vol. 83, No. 4, pp. 87-96, 2015.
- [8] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning About Knowledge*. MIT Press, Cambridge, MA, 2003.
- [9] J. Halpern and R. Pass. Conservative belief and rationality. *Games and Economic Behavior*, Vol. 80, pp. 186-192, 2013.
- [10] J. Harsanyi. Games with Incomplete Information Played by “Bayesian” Players, I-III. *Management Science*, Vol. 14, pp. 159-182, 320-334, 486-502. 1967-1968.
- [11] S. Kripke. Semantical analysis of modal logic I: normal modal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, Vol. 9, pp. 67-96, 1963.
- [12] M. Siniscalchi. Epistemic Game Theory: Beliefs and Types. In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*, Second Edition. Palgrave Macmillan, New York, 2008.
- [13] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, Vol. 16, No. 1, pages 8-37, Mar., 1961.